arXiv:1702.07605v1 [cs.DC] 24 Feb 2017

Compact Self-Stabilizing Leader Election for Arbitrary Networks

Lélia Blin* Sébastien Tixeuil†

Abstract

We present a self-stabilizing leader election algorithm for arbitrary networks, with space-complexity $O(\max\{\log \Delta, \log\log n\})$ bits per node in n-node networks with maximum degree Δ. This space complexity is sub-logarithmic in n as long as $\Delta = n^{o(1)}$. The best space-complexity known so far for arbitrary networks was $O(\log n)$ bits per node, and algorithms with sub-logarithmic space-complexities were known for the ring only. To our knowledge, our algorithm is the first algorithm for self-stabilizing leader election to break the $\Omega(\log n)$ bound for silent algorithms in arbitrary networks. Breaking this bound was obtained via the design of a (non-silent) self-stabilizing algorithm using sophisticated tools such as solving the distance-2 coloring problem in a silent self-stabilizing manner, with space-complexity $O(\max\{\log \Delta, \log\log n\})$ bits per node. Solving this latter coloring problem allows us to implement a sub-logarithmic encoding of spanning trees — storing the IDs of the neighbors requires $\Omega(\log n)$ bits per node, while we encode spanning trees using $O(\max\{\log \Delta, \log\log n\})$ bits per node. Moreover, we show how to construct such compactly encoded spanning trees without relying on variables encoding distances or number of nodes, as these two types of variables would also require $\Omega(\log n)$ bits per node.

*Sorbonne Universités, UPMC Univ Paris 06, CNRS, Université d'Evry-Val-d'Essonne, LIP6 UMR 7606, 4 place Jussieu 75005, Paris.

†Sorbonne Universités, UPMC Univ Paris 06, CNRS, LIP6 UMR 7606, 4 place Jussieu 75005, Paris.

1 Introduction

This paper tackles the problem of designing memory efficient self-stabilizing algorithms for the leader election problem. *Self-stabilization* [17, 18, 33] is a general paradigm to provide recovery capabilities to networks. Intuitively, a protocol is self-stabilizing if it is able to recover from any transient failure, without external intervention. *Leader election* is one of the fundamental building blocks of distributed computing, as it enables a single node in the network to be distinguished, and thus to perform specific actions. Leader election is especially important in the context of self-stabilization as many protocols for various problems assume that a single leader exists in the network, even after faults occur. Hence, a self-stabilizing leader election mechanism enables such protocols to be run in networks where no leader is given a priori, by using simple stabilization-preserving composition techniques [18]. *Memory efficiency* relates to the amount of information to be sent to neighboring nodes for enabling stabilization. A small space-complexity induces a smaller amount of information transmission, which (1) reduces the overhead of self-stabilization when there are no faults, or after stabilization [1], and (2) facilitates mixing self-stabilization and replication [24, 26].

A foundational result regarding space-complexity in the context of self-stabilizing silent algorithms[1] is due to Dolev et al. [19], stating that, in n-node networks, $\Omega(\log n)$ bits of memory per node are required for solving tasks such as leader election. So, only *talkative* algorithms may have $o(\log n)$-bit space-complexity for self-stabilizing leader election. Several attempts to design compact self-stabilizing leader election algorithms (i.e., algorithms with space-complexity $o(\log n)$ bits) were performed, but restricted to rings. The algorithms by Mayer et al. [32], by Itkis and Levin [28], and by Awerbuch and Ostrovsky [5] use a constant number of bits per node, but they only guarantee *probabilistic* self-stabilization (in the Las Vegas sense). *Deterministic* self-stabilizing leader election algorithms for rings were first proposed by Itkis et al. [29] for rings with a prime number of nodes. Beauquier et al. [6] consider rings of arbitrary size, but assume that node identifiers in n-node rings are bounded from above by $n + k$, where k is a small constant. A recent result by Blin et al. [13] demonstrates that both previous constraints in the deterministic setting are unnecessary, by presenting a deterministic self-stabilizing leader election algorithm for rings of arbitrary size using identifiers of arbitrary polynomially bounded values, with space complexity $O(\log \log n)$ bits.

In general networks, self-stabilizing leader election is tightly connected to self-stabilizing tree-construction. On the one hand, the existence of a leader permits time- and memory-efficient self-stabilizing tree-construction [14, 20, 15, 11, 31]. On the other hand, growing and merging trees is the main technique for designing self-stabilizing leader election algorithms in networks, as the leader is often the root of an inward tree [3, 4, 2, 9]. To the best of our knowledge, all algorithms that do not assume a pre-existing leader [3, 4, 2, 8] for tree-construction use $\Omega(\log n)$ bits per node. This high space-complexity is due to the implementation of two main techniques, used by all algorithms, and recalled below.

The first technique is the use of a *pointer-to-neighbor* variable, that is meant to designate unambiguously one particular neighbor of every node. For the purpose of tree-construction, pointer-to-neighbor variables are typically used to store the parent node in the constructed tree. Specifically, the parent of every node is designated unambiguously by its identifier, requiring $\Omega(\log n)$ bits for each pointer variable. In principle, it would be possible to reduce the memory to $O(\log \Delta)$ bits per pointer variable in networks with maximum degree Δ, by using node-coloring at distance 2 instead of identifiers to identify neighbors. However, this in turn would require the availability of a self-stabilizing distance-2 node-coloring algorithm that uses $o(\log n)$ bits per node. Previous

[1] An algorithm is *silent* if each of its executions reaches a point in time after which the states of nodes do not change. A non-silent algorithm is said to be *talkative* (see [13]).

self-stabilizing distance-2 coloring algorithms use variables of large size. For instance, in the algorithm by Herman et al. [27], every node communicates its distance-3 neighborhood to all its neighbors, which yields a space-complexity of $O(\Delta^3 \log n)$ bits. Johnen et al. [25] draw random colors in the range $[0, n^2]$, which yields a space-complexity of $O(\log n)$ bits. Finally, while the deterministic algorithm of Blair et al. [7] reduces the space-complexity to $O(\log \Delta)$ bits per node, this is achieved by ignoring the cost for storing another pointer-to-neighbor variable at each node. In absence of a distance-2 coloring (which their algorithm [7] is precisely supposed to produce), their implementation still requires $\Omega(\log n)$ bits per node. To date, no self-stabilizing algorithm implement pointer-to-neighbor variables with space-complexity $o(\log n)$ bits in arbitrary networks.

The second technique for tree-construction or leader election is the use of a *distance* variable that is meant to store the distance of every node to the elected node in the network. Such distance variable is used in self-stabilizing spanning tree-construction for breaking cycles resulting from arbitrary initial state (see [3, 4, 2]). Clearly, storing distances in n-node networks may require $\Omega(\log n)$ bits per node. There are a few self-stabilizing tree-construction algorithms that are not using explicit distance variables (see, e.g., [30, 22, 16]), but their space-complexity is $O(n \log n)$ bits [22, 16] or $O(\log n + \Delta)$ [30]. Using the general principle of distance variables with space-complexity below $\Theta(\log n)$ bits was attempted by Awerbuch et al. [5], and Blin et al. [12]. These papers distribute pieces of information about the distances to the leader among the nodes according to different mechanisms, enabling to store $o(\log n)$ bits per node. However, these sophisticated mechanisms have only been demonstrated in rings. To date, no self-stabilizing algorithms implement distance variables with space-complexity $o(\log n)$ bits in arbitrary networks.

Our results

In this paper, we design and analyze a self-stabilizing leader election algorithm with space-complexity $O(\max\{\log \Delta, \log \log n\})$ bits in n-node networks with maximum degree Δ. This algorithm is the first self-stabilizing leader election algorithm for arbitrary networks with space-complexity $o(\log n)$ (whenever $\Delta = n^{o(1)}$). It is designed for the standard state model (a.k.a. shared memory model) for self-stabilizing algorithms in networks, and it performs against the unfair distributed scheduler.

The design of our algorithm requires to overcome several bottlenecks, including the difficulties of manipulating pointer-to-neighbor and distance variables using $o(\log n)$ bits in arbitrary networks. Overcoming these bottlenecks was achieved thanks to the development of sub-routine algorithms, each deserving independent special interest, described hereafter.

First, we generalize to arbitrary networks the techniques proposed [12, 13] for rings, and aiming at publishing the identifiers in a bit-wise manner. This generalization allows us to manipulate the identifiers with just $O(\log \log n)$ bits of memory per node.

Second, we propose a *silent* self-stabilizing algorithm for distance-2 coloring with space-complexity $O(\max\{\log \Delta, \log \log n\})$ bits. As opposed to previous distance-2 coloring algorithms, we do not use identifiers for encoding pointer-to-neighbor variables, but we use a compact representation of the identifiers to break symmetries. This allows us to design a compact encoding of spanning trees.

Third, we design a new technique to detect the presence of cycles in the initial configuration resulting from a transient failure. This technique does not use distances, but is based on the uniqueness of each identifier in the network. Notably, this technique can be implemented by a *silent* self-stabilizing algorithm, with space-complexity $O(\max\{\log \Delta, \log \log n\})$ bits per node.

Last but not least, we design a new technique to avoid the creation of cycles during the execution of the leader election algorithm. Again, this technique does not uses distances but maintains a spanning forest, which eventually reduces to a single spanning tree rooted at the leader at the completion of the leader election algorithm. Implementing this technique results in a self-stabilizing algorithm with space complexity $O(\max\{\log \Delta, \log \log n\})$ bits per node.

2 Model and definitions

2.1 Protocol syntax and semantics

We consider a distributed system consisting of n processes that form a arbitarry communication graph. The processes are represented by the nodes of this graph, and the edges represent pairs of processes that can communicate directly with each other. Such processes are said to be *neighbors*. Let $G = (V, E)$ be an n-node graph, where V is the set of nodes, and E the set of edges and Δ the degree of the graph. A node v has access to a constant unique identifier ID_v, but can only access its identifier one bit at a time, using the $\text{Bit}(x, \text{ID}_v)$ function, which returns the position of the x^{th} most significant bit equal to 1 in ID_v. This position can be encoded with $O(\log \log n)$ bits when identifiers are encoded using $O(\log n)$ bits, as we assume they are. A node v has access to locally unique port numbers associated to its adjacent edges. We do not assume any consistency between port numbers of a given edge. In short, port numbers are constant throughout the execution but initialized by an adversary. Each process contains variables and rules. A variable ranges over a domain of values. The variable var_v denote the variable var located at node v. A rule is of the form $\langle label \rangle : \langle guard \rangle \longrightarrow \langle command \rangle$. A *guard* is a boolean predicate over process variables. A *command* is a set of variable-assignments. A command of process p can only update its own variables. On the other hand, p can read the variables of its neighbors. This classical communication model is called the *state model* or the *state-sharing communication model*.

An assignment of values to all variables in the system is called a *configuration*. A rule whose guard is **true** in some system configuration is said to be *enabled* in this configuration. The rule is *disabled* otherwise. The atomic execution of a subset of enabled rules (at most one rule per process) results in a transition of the system from one configuration to another. This transition is called a *step*. A *run* of a distributed system is a maximal alternating sequence of configurations and steps. Maximality means that the execution is either infinite, or its final configuration has no rule enabled.

2.2 Schedulers

The asynchronism of the system is modeled by an adversary (a.k.a. *scheduler*) that chooses, at each step, the subset of enabled processes that are allowed to execute one of their rules during this step. The literature proposed a lot of daemons depending of their characteristics (like fairness, distribution, ...), see [21] for a taxonomy of these scheduler. Note that we assume here an unfair distributed scheduler. This scheduler is the most challenging since no assumption is made of the subset of enabled processes chosen by the scheduler at each step (That only require this set to be non empty if the set of enabled processes is not empty in order to guarantee progress of the algorithm.)

2.3 Predicates and specifications

A predicate is a boolean function over configurations. A configuration *conforms* to some predicate R, if R evaluates to **true** in this configuration. The configuration *violates* the predicate otherwise. Predicate R is *closed* in a certain protocol P, if every configuration of a run of P conforms to R, provided that the protocol starts from a configuration conforming to R. Note that if a protocol configuration conforms to R, and the configuration resulting from the execution of any step of P also conforms to R, then R is closed in P.

Problem specification prescribes the protocol behavior. The output of the protocol is carried through external variables, that are updated by the protocol, and used to display the results of the protocol computation. The problem specification is the set of sequences of configurations of external variables.

A protocol implements the specification. Part of the implementation is the mapping from the protocol configurations to the specification configurations. This mapping does not have to be one-to-one. However, we only consider unambiguous protocols where each protocol configuration maps to only one specification configuration. Once the mapping between protocol and specification configurations is established, the protocol runs are mapped to specification sequences as follows. Each protocol configuration is mapped to the corresponding specification configuration. Then, stuttering, the consequent identical specification configurations, is eliminated. Overall, a run of the protocol satisfies the specification if its mapping belongs to the specification. Protocol P *solves* problem S under a certain scheduler if every run of P produced by that scheduler satisfies the specifications defined by S. Given two predicates l_1 and l_2 for protocol P, l_2 is an *attractor* for l_1 if every run that starts from a configuration that conforms to l_1 contains a configuration that conforms to l_2. Such a relationship is denoted by $l_1 \rhd l_2$. Also, the $\rhd$ relation is transitive: if l_1, l_2, and l_3 are predicates for P, and $l_1 \rhd l_2$ and $l_2 \rhd l_3$, then $l_1 \rhd l_3$. In this last case, l_2 is called an *intermediate* attractor towards l_3.

Definition 1 (Self-stabilization)**.** *A protocol P is* self-stabilizing *[17] to specification S if there exists a predicate L for P such that:*

1. *L is an attractor for* true*,*
2. *Any run of P starting from a configuration satisfying L satisfies S.*

Definition 2 (Leader Election)**.** *Consider a system of processes where each process' set of variables is mapped to a boolean specification variable leader denoted by ℓ. The* leader election *specification sequence consists in a single specification configuration where a unique process p maps to $\ell_p = true$, and every other process $q \neq p$ maps to $\ell_q = false$.*

3 Compact self-stabilizing leader election for networks

Our self-stabilizing leader election algorithm is based on a spanning tree-construction rooted at a maximum degree node, without using distances. If multiple maximum degree nodes are present in the network, we break ties with colors and if necessary with identifiers.

Theorem 1. *Algorithm* **C-LE** *solves the leader election problem in a talkative self-stabilizing manner in any n-node graph, assuming the state model and a distributed unfair scheduler, with $O(\max\{\log \Delta, \log\log n\})$ bits of memory per node, where Δ is the graph's degree.*

Our talkative self-stabilizing algorithm reuses and extends a technique for obtaining compact identifiers of size $O(\log\log n)$ bits per node presented in Section 3.1. Then, the leader election process consists in running several algorithms layers using decreasing priorities:

1. A silent self-stabilizing distance-2 coloring presented in subsection 3.2 that permits to implement pointer-to-neighbors with $o(\log n)$ bits per node.
2. A silent self-stabilizing cycle and illegitimate sub spanning tree destruction reused from previous work [10, 13] presented in subsection 3.3.
3. A silent self-stabilizing cycle detection that does not use distance to the root variables presented in subsection 3.4.
4. A talkative self-stabilizing spanning tree-construction, that still does not use distance to the root variables, presented in subsection 3.5. This algorithm is trivially modified to obtain a leader election algorithm.

Due to the lack of space most of the proofs and predicates are delegated to the Appendix.

3.1 Compact memory using identifiers

As many deterministic self-stabilizing leader election algorithms, our approach ends up comparing node unique identifiers. However, to avoid communicating the full $\Omega(\log n)$ bits to each neighbor at any given time, we reuse the scheme devised in previous work [12, 13] to progressively publish node identifiers. Let ID_v be the identifier of node v. We assume that $\text{ID}_v = \sum_{i=0}^{k} b_i 2^i$. Let $I_v = \{i \in \{0, ..., k\}, b_i \neq 0\}$ be the set of all non-zero bit-positions in the binary representation of ID_v. Then, I_v can be written as $\{pos_1, ..., pos_j\}$, where $pos_k > pos_{k+1}$. In the process of comparing node unique identifiers during the leader election algorithm execution, the nodes must first agree on the same bit-position pos_{j-i+1} (for $i = 1, \ldots, j$); this step of the algorithm defines *phase* i. Put differently, the bit-positions are communicated in decreasing order of significance in the encoding of the identifier.

If all identifiers are in $[1, n^c]$, for some constant $c \geqslant 1$, then the communicated bit-positions are less than or equal to $c\lceil \log n \rceil$, and thus can be represented with $O(\log \log n)$ bits. However, the number of bits used to encode identifiers may be different for two given nodes, so there is no common upper bound for the size of identifiers. Instead, we use variable $\widehat{\mathsf{B}}_v$, which represents the most significant bit-position of node v. In other words, $\widehat{\mathsf{B}}_v$ represents the size of the binary representation of ID_v. The variables ph, Bp are the core of the identifier comparison process. Variable ph_v stores the current phase number i, while Variable Bp_v stores the bit-position of ID_v at phase i. Remark that the number of non-zero bits can be smaller than the size of the binary representation of the identifier of the node, so if there are no more non-zero bit at phase $i \leqslant \widehat{\mathsf{B}}_v$, we use $\mathsf{Bp}_v = -1$. To make the algorithm more readable, we introduce Variable $\mathcal{C}id_v = (\widehat{\mathsf{B}}_v, \mathsf{ph}_v, \mathsf{Bp}_v)$, called a *compact identifier* in the sequel. When meaningful, we use $\mathcal{C}id_v^i = (\widehat{\mathsf{B}}_v, \mathsf{Bp}_v)$, where $i = \mathsf{ph}_v$.

Node v can trivially detect an error (see predicate $\mathtt{ErT}(v)$) whenever its compact identifier does not match its global identifier, or its phase is greater than $\widehat{\mathsf{B}}_v$. Moreover, the phases of neighboring nodes must be close enough: a node's phase may not be more than 1 ahead or behind any of its neighbors; also a node may not have a neighbor ahead and another behind. Predicate $\mathtt{SErB}(v, S)$ captures these conditions, where $S(v)$ denotes a subset of neighbors of v. The set S should be understood as an input provided by an upper layer algorithm. If v detects an error through $\mathtt{ErT}(v)$ or $\mathtt{SErB}(v, S)$, it resets its compact identifier to its first phase value (see command $\mathcal{R}esetCid(v)$). The compact identifier of u is smaller (respectively greater) than the compact identifier of v, if the most significant bit-position of u is smaller (respectively greater) than the most significant bit-position of v, or if the most significant bit-position of u is equal to the most significant bit-position of v, u and v are in the same phase, and the bit-position of u is smaller (respectively greater) than the bit-position of v:

$$\mathcal{C}id_u^i <_c \mathcal{C}id_v^i \equiv (\widehat{\mathsf{B}}_u < \widehat{\mathsf{B}}_v) \vee \big((\widehat{\mathsf{B}}_v = \widehat{\mathsf{B}}_u) \wedge (\mathsf{Bp}_u < \mathsf{Bp}_v)\big) \tag{1}$$

When two nodes u and v have the same most significant bit-position and the same bit position at phase $i < \widehat{\mathsf{B}}_v$, they are possibly equal with respect to compact identifiers (denoted by $\simeq_c$).

$$\mathcal{C}id_u^i \simeq_c \mathcal{C}id_v^i \equiv (i < \widehat{\mathsf{B}}_v) \wedge \big((\widehat{\mathsf{B}}_v = \widehat{\mathsf{B}}_u) \wedge (\mathsf{Bp}_u = \mathsf{Bp}_v)\big) \tag{2}$$

Finally, two nodes u and v have the same compact identifier (denoted by $=_c$) if their phase reaches the size of the binary representation of the identifier of the two nodes, and their last bit-position is the same.

$$\mathcal{C}id_u =_c \mathcal{C}id_v \equiv (i = \widehat{\mathsf{B}}_v = \widehat{\mathsf{B}}_u) \wedge (\mathsf{Bp}_u = \mathsf{Bp}_v) \tag{3}$$

The predicates $\mathtt{SPh}^+(v)$ and $\mathtt{TPh}^+(v, S)$ check if a node v can increases its phase (or restarts $\mathcal{C}id_v$), the first one is dedicated to the silent protocols, the second one is dedicated to the talkative protocols. The command $\mathcal{I}ncPh(v)$ is dedicated for increasing phases or restarting $\mathcal{C}id_v$. Last, the command $\mathcal{O}pt$ assigns at a node v the minimum (or maximum) compact identifier in the subset of neighbors $S(v)$. We have now, all the principals ingredients to use compact identifiers.

3.2 Silent self-stabilizing distance-2 coloring

In this section, we provide a solution to assign colors that are unique up to distance two (and bounded by a polynom of the graph degree) in any graph. Those colors are meant to efficiently implement the pointer-to-neighbor mechanism that otherwise requires $\Omega(\log n)$ bits per node.

Our solution uses compact identifiers to reduce memory usage. When a node v has the same color as (at least one of) its neighbors, then if the node v has the smallest conflicting color in its neighborhood and is not the biggest identifier among conflicting nodes, then v changes its color. To make sure a fresh color is chosen by v, all nodes publish the maximum color used by their neighborhood (including themself). So, when v changes its color, it takes the maximum advertised color plus one. Conflicts at distance two are resolved as follow: let us consider two nodes u and v in conflict at distance two, and let w be (one of) their common neighbor; as w publishes the color of u and v, it also plays the role of a relay, that is, w computes and advertises the maximum identifiers between u and v, using the compact identifiers mechanisms that were presented above; a bit by bit, then, if v has the smallest identifier, it changes its color to a fresh one. To avoid using too many colors when selecting a fresh one, all changes of colors are made modulo an upper bound on the number of neighbors at distance 2, which is computed locally by each node.

3.2.1 Self-stabilizing algorithm description

Each node v maintains a color variable denoted by $\mathtt{c}_v$ and a degree variable denoted by δ_v. A variable $\check{c}_v$ stores the minimum color in conflict in its neighborhood (including itself). The variable $\hat{c}_v$ stores the maximum color observed in its neighborhood. We call v a *player* node when v has the minimum color in conflict. Also, we call u a *relay* node when u does not have the minimum color in conflict, yet at least two of its neighbors have the minimum color in conflict.

The rule $\mathbb{R}_{\Delta}$ assures that the degree variable is equal to the degree of the node. Each node v must maintain its color in range $[1, \Delta(v)^2 + 1]$ to satisfy the memory requirements of our protocol, where $\Delta(v)$ is a function that returns the maximum degree of its neighborhood (including itself). Whenever v's color exceeds its expected range, rule $\mathbb{R}^+_{\Delta}$ resets the color to one. Rule $\mathbb{R}_{\mathtt{Up}}$ is dedicated to updating the variables of v whenever they do not match the observed neighborhood of v (see $\mathtt{Bad}(v)$), or when a player node has an erroneous phase variable when comparing its identifier with another player node (see function $\mathtt{Oth}(v)$). In both cases, the v computes the minimum and maximum color and resets its compact identifier variable (see command $\mathcal{U}pdate(v)$). The rule $\mathbb{R}_{\mathtt{Color}}$ increases the color of the node v but maintains the color in some range (see command $\mathcal{N}ewcolor(v)$), when v has the minimum color in conflict and the minimum identifier. The rule $\mathbb{R}_{\mathtt{Bit}}$ increases the phase of the node v, when v is a player and does not have the minimum identifier at the selected phase. The rule $\mathbb{R}_{\mathtt{Relay}}$ updates the identifier variable when v is a relay node.

Algorithm 1: C-Color

$$
\begin{array}{lll}
\mathbb{R}_{\Delta} & : (\delta_v \neq \deg(v)) & \longrightarrow \quad \delta_v := \deg(v); \\
\mathbb{R}^+_{\Delta} & : (\delta_v = \deg(v)) \wedge (\mathtt{c}_v > \Delta(v)^2) & \longrightarrow \quad \mathtt{c}_v := 1; \\
\mathbb{R}_{\mathtt{Up}} & : (\delta_v = \deg(v)) \wedge (\mathtt{c}_v \leqslant \Delta(v)^2) \wedge \mathtt{Bad}(v) & \longrightarrow \quad \mathcal{U}pdate(v); \\
\mathbb{R}_{\mathtt{Color}} & : (\delta_v = \deg(v)) \wedge (\mathtt{c}_v \leqslant \Delta(v)^2) \wedge \neg\mathtt{Bad}(v) \wedge \mathtt{Player}(v) \wedge \mathtt{Loser}(v) & \longrightarrow \quad \mathcal{N}ewcolor(v); \\
\mathbb{R}_{\mathtt{Bit}} & : (\delta_v = \deg(v)) \wedge (\mathtt{c}_v \leqslant \Delta(v)^2) \wedge \neg\mathtt{Bad}(v) \wedge \mathtt{Player}(v) \wedge \neg\mathtt{Loser}(v) \wedge \mathtt{SPh}^+(v, \mathtt{Oth}(v)) & \longrightarrow \quad \mathcal{I}ncPh(v); \\
\mathbb{R}_{\mathtt{Relay}} & : (\delta_v = \deg(v)) \wedge (\mathtt{c}_v \leqslant \Delta(v)^2) \wedge \neg\mathtt{Bad}(v) \wedge \mathtt{Relay}(v) \wedge \mathtt{RUp}(v, \mathtt{PlayR}(v)) & \longrightarrow \quad \mathcal{O}pt(v, \mathtt{PlayR}(v), \max);
\end{array}
$$

Theorem 2. *Algorithm* **C-Color** *solves the vertex coloration problem at distance two in a silent self-stabilizing manner in graph, assuming the state model, and a distributed unfair scheduler. Moreover, if the n node identifiers are in $[1, n^c]$, for some $c \geqslant 1$, then Algorithm* **C-Color** *uses $O(\max\{\log \Delta, \log\log n\})$ bits of memory per node.*

3.3 Cleaning a cycle or an impostor-rooted spanning tree

The graph G is supposed to be colored up to distance 2, thanks to our previous algorithm. To construct a spanning tree of G, each node v maintains a variable p_v storing the color of v's parent ($\varnothing$ otherwise). The function $\mathtt{Ch}(v)$ to return the subset of v's neighbors considered as its children (that is, each such node u has its p_u variable equal v's color). Not that the variable parents is managed by the algorithm of spanning tree-construction.

An error is characterized by the presence of inconsistencies between the values of the variables of a node v and those of its neighbors. In the process of a tree-construction, an error occurring at node v may have impact on its descendants. For this reasons, after a node v detects an error, our algorithm cleans v and all of its descendants. The cleaning process is achieved by Algorithm **Freeze**, already presented in previous works [12, 13, 10]. Algorithm **Freeze** is run in two cases: cycle detection (thanks to predicate ErCycle(v), presented in Subsection 3.4), and impostor leader detection (thanks to predicate ErST(v), presented in Subsection 3.5). An impostor leader is a node that (erroneously) believes that it is a root.

When a node v detects a cycle or an impostor root, v deletes its parent. Simultaneously, v becomes a *frozen* node. Then, every descendant of v becomes frozen. Finally, from the leaves of the spanning tree rooted at v, nodes delete their parent and reset all variables that are related to cycle detection or tree-construction. So, this cleaning processe cannot create a livelock. Algorithm **Freeze** is a silent self-stabilizing algorithm using $O(1)$ bits of memory per node (see Annexe A.3).

3.4 Silent self-stabilizing algorithm for cycle detection

We present in this subsection a self-stabilizing algorithm to detect cycles (possibly due to initial incorrect configuration) without using the classical method of computing the distance to the root. We first present our solution with the assumption of global identifiers (hence using $O(\log n)$ bits for an n-node network), and then using our compact identifier scheme.

3.4.1 Self-stabilizing algorithm with identifiers

The main idea to detect cycles is to use the uniqueness of the identifiers. We flow the minimum identifier up the tree to the root, then if a node whose identifier is minimum receives its identifier, it can detect a cycle. Similarly, if a node v has two children flowing the same minimum identifier, v can detect a cycle. The main issue to resolve is when the minimum identifier that is propagated to the root does not exists in the network (that is, it results from an erroneous initial state).

The variable m_v stores the minimum identifier collected from the leaves to the root up to node v. We denote by E_v the minimum identifier obtained by v during the previous iteration of the protocol (this can be $\varnothing$). A node v may selects among its children the node u with the smallest propagated identifier stored in m_u, we call this child *kid* returned by the function $\mathsf{k}(v)$.

Predicate ErCycle(v) is the core of our algorithm. Indeed, a node v can detect the presence of a cycle if it has a parent and if *(i)* one of its child publishes its own identifier, or *(ii)* two of its children publish the same identifier. Let us explain those conditions in more detail.

Let us consider a spanning structure S, a node $v \in S$ and let u and w be two of its children. Suppose that v and u belong to a cycle $\mathtt{C}$, note that, since a node has a single parent, w cannot belong to any cycle (see Figure 1). Let $\breve{\mathsf{m}}$ be the minimum identifier stored by any variable m_z such that z belongs to S. So, z is either in $\mathtt{C}$, or in the subtree rooted to w, denoted by $\mathtt{T}_w$.

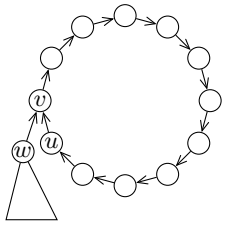

Fig. 1 – Spanning structure

First, let us consider the case where $\breve{\mathsf{m}}$ is stored in $\mathtt{T}_w$. As any node selects the minimum for flowing the m upstream, there exists a configu-

ration γ where $\mathsf{m}_w = \breve{\mathsf{m}}$, and a configuration $\gamma' > \gamma$ where $\mathsf{m}_u = \breve{\mathsf{m}}$. In γ', v can detect an error, due to the uniqueness of identifier, it is not possible for two children of v to share the same value when there is no cycle.

Now, let us suppose that $\breve{\mathsf{m}}$ is in $\mathtt{C}$, and let v' be the node with the smallest identifier in $\mathtt{C}$, so $\mathsf{m}_{v'} = \breve{\mathsf{m}}$ or $\mathsf{m}_{v'} \neq \breve{\mathsf{m}}$ ($\mathsf{m}_{v'} \neq \breve{\mathsf{m}}$ means that the identifier $\breve{\mathsf{m}}$ does not exist in $\mathtt{C}$.) If $\mathsf{m}_{v'} = \breve{\mathsf{m}}$, as any node selects the minimum for flowing the m upstream, there exists a configuration γ where $\mathsf{m}_{u'} = \breve{\mathsf{m}}$ and u' is the child of v' involved in $\mathtt{C}$, then v' can detect an error. Indeed, due to the uniqueness of identifier, it is not possible that one of its children store its identifier when there is no cycle. The remaining case is when $\mathsf{m}_{v'} \neq \breve{\mathsf{m}}$. In this case, as any node selects the minimum for flowing the m upstream, there exists a configuration γ where $\mathsf{m}_z = \breve{\mathsf{m}}$, with z belonging to $\mathtt{C}$. When a node v, its parent and one of its children share the same minimum, they restart the computation of the minimum identifier. For this purpose, they put their own identifier in the m variable. To avoid livelock, they also keep track of the previous $\breve{\mathsf{m}}$ in variable E_v. Now $\breve{\mathsf{m}} = \mathsf{m}_{v'}$, so the system reaches the first case. Note that the variable E_v blocks the livelock but also the perpetual restart of the nodes, as a result of this, a silent algorithm. Moreover, a node v collects the minimum identifier from the leaves to the root, if m_v contains an identifier bigger than the identifier of the node v, then v detects an error. The same holds, when v has a m_v smaller than m_u with u children of v, since the minimum is computed between $\mathsf{m}_{\mathsf{k}(v)}$ and its own identifier.

$$\mathrm{ErCycle}(v) \equiv (\mathsf{p}_v \neq \varnothing) \wedge \Big((\mathsf{m}_{\mathsf{k}(v)} = \mathrm{ID}_v) \vee (\exists (u,w) \in \mathtt{Ch}(v) : \mathsf{m}_u = \mathsf{m}_w) \vee (\mathsf{m}_v > \mathrm{ID}_v) \vee \big((\mathsf{m}_v \neq \mathrm{ID}_v) \wedge (\mathsf{m}_v < \mathsf{m}_{\mathsf{k}(v)})\big)\Big)$$

Our algorithm contains three rules. The first rule $\mathbb{R}_{\mathtt{Min}}(v)$ updates the minimum variable m_v if the minimum variable m_u of a child u is smaller, nevertheless this rule is enabled if and only if the variable E_v does not contain the minimum m_u published by the child. When a node v and its relatives have the same minimum, v declares its intent to restart a minimum identifier computation by erasing its current (and storing it in E_v). The rule $\mathbb{R}_{\mathtt{Start}}(v)$ is dedicated to declaring its intent to restart. When all its neighbors have the same intent, the node can restart (see rule $\mathbb{R}_{\mathtt{ID}}(v)$).

Algorithm 2: Algorithm **Break** For node v with $\neg\mathrm{ErCycle}(v)$

$$\begin{array}{lll}
\mathbb{R}_{\mathtt{Min}} & : (\mathsf{m}_v > \mathsf{m}_{\mathsf{k}(v)}) \wedge (\mathsf{E}_v \neq \mathsf{m}_{\mathsf{k}(v)}) & \longrightarrow \mathsf{m}_v := \mathsf{m}_{\mathsf{k}(v)}; \\
\mathbb{R}_{\mathtt{Start}} & : (\mathsf{m}_{\mathsf{p}_v} = \mathsf{m}_v = \mathsf{m}_{\mathsf{k}(v)}) \wedge (\mathsf{E}_v \neq \mathsf{m}_v) & \longrightarrow \mathsf{E}_v := \mathsf{m}_v; \\
\mathbb{R}_{\mathtt{ID}} & : (\mathsf{E}_{\mathsf{p}_v} = \mathsf{E}_v = \mathsf{E}_{\mathsf{k}(v)} = \mathsf{m}_v) \wedge (\mathsf{m}_v \neq \mathrm{ID}_v) & \longrightarrow \mathsf{m}_v := \mathrm{ID}_v;
\end{array}$$

Theorem 3. *Algorithm* **Break** *solves the detection of cycle in n-node graph in a silent self-stabilizing manner, assuming the state model, and a distributed unfair scheduler. Moreover, if the n node identifiers are in $[1, n^c]$, for some $c \geqslant 1$, then algorithm* **Break** *uses $O(\log n)$ bits of memory per node.*

3.4.2 Talkative self-stabilizing cycle detection with compact identifiers

We refine algorithm **Break** to make use of compact identifiers (of size $O(\log\log n)$ instead of global identifiers (of size $O(\log n)$). With compact identifiers, the main problem is the following: two nodes u and v can deduce that $\mathcal{C}id_u =_c \mathcal{C}id_u$ if and only if they have observed $\mathcal{C}id_u \simeq_c \mathcal{C}id_v$ during every phase i, with $1 \leqslant i \leqslant \widehat{\mathsf{B}}_v$. A node v selects the minimum compact identifier stored in variable m in its neighborhood (including itself). If in a previous configuration one of its children had presented v a compact identifier smaller than its own, v became passive (Variable $\mathtt{Active}_v = \mathit{false}$), and remained active otherwise (Variable $\mathtt{Active}_v = \mathit{true}$). Only active nodes can continue to increase their phase. Moreover, a node increases its phase if and only if its parent and one of its children u has the same information, namely $\mathcal{C}id_u \simeq_c \mathcal{C}id_v \simeq_c \mathcal{C}id_{\mathsf{p}_v}$. Note that, in a spanning tree several nodes may not increase their phases, for example leaves which by definition have no child, this does

not cause a problem. Let us explain, let v be the node with the smallest identifier involved in a cycle and let suppose that v has two children, one u involved in the cycle and the child w no. In some configuration the node w has not able to increase its phase, but the node u will reach the same phase of the active node v, so v increase its phase, and the system reaches a configuration where $\mathcal{C}id_u =_c \mathcal{C}id_v$ so v detects an error of cycle. The variable E_v combined to this compact identifier use allow us to maintain a silent algorithm.

Predicate ErCycle now takes into account the error(s) related to compact identifiers management. It is important to note that the cycle breaking algorithm does not manage phase differences. Indeed, a node v with a phase bigger than the phase of one of its children u takes the m_u, if and only if its phase its bigger than two or if no child has its same compact identifier. The m_v variable is be compared using lexicographic order by rule $\mathbb{R}_{\texttt{Min}}$. Tthe modifications to algorithm **Break** are minor. We add only one rule to increases the phase: $\mathbb{R}_{\texttt{Inc}}$. Only a passive node can restart. Remark that now the m variable uses $O(\log\log n)$ bits. As the p variables stores a color, we obtain a memory requirement of $O(\max\{\log\Delta, \log\log n\})$ bits per node.

Theorem 4. *Algorithm* **C-Break** *solves the detection of cycle in arbitrary n-node graph in a silent self-stabilizing manner, assuming the state model, and a distributed unfair scheduler. Moreover, if the n node identifiers are in $[1, n^c]$, for some $c \geqslant 1$, then algorithm* **C-Break** *uses $O(\max\{\log\Delta, \log\log n\})$ bits of memory per node.*

3.5 Spanning tree-construction without distance to the root maintenance

Our approach for self-stabilizing leader election is to construct a spanning tree whose root is to be the elected leader. Two main obstacles to self-stabilizing tree-construction are the possibility of an arbitrary initial configuration containing one or more cycles, or the presence of one or more impostor-rooted spanning trees. We already explained how the cycle detection and cleaning process takes place, so we focus in this section on *cycleless* configurations.

The main idea is to mimics the *fragments* approach introduced by Gallager *et al.* [23]. In an ideal situation, at the beginning each node is a fragment, each fragment merges with a neighbor fragment holding a bigger root signature, and at the end remains only one fragment, rooted in the root with the biggest signature (that is, the root with maximum degree, maximum color, and maximum global identifier). To maintain a spanning structure, the neighbors that become relatives (that is, parents or children) remain relatives thereafter. Note that the relationship may evolve through time (that is, a parent can become a child and vice versa). So our algorithm maintains that as an invariant, given by Lemma 10.

Indeed, when two fragments merge, the one with the root with smaller signature F_1 and the other one with a root with bigger signature F_2, the root of F_1 is re-rooted toward its descendants until reaching the node that identified F_2. This approach permits to construct an acyclic spanning structure, without having to maintain distance information. The variable R_v stores the signature relative to the root (that is, its degree, its color, and its identifier). Note that, the comparison between two R is done using lexical ordering. Moreover, the variable new_v stores the color of the neighbor w of v leading to the a node u with $\mathsf{R}_u > \mathsf{R}_v$ if there exists such a node, and $\varnothing$ otherwise. The function $f(v)$ returns the color of the neighbor of v with the maximum root.

Let us now give more details about our algorithm (presented in Algorithm 3). If a root v has a neighbor u with $\mathsf{R}_u > \mathsf{R}_v$, then v chooses u as its parent (see rule $\mathbb{R}_{\texttt{Merge}}$). If a node v (not a root) has a neighbor u with $\mathsf{R}_u > \mathsf{R}_v$, it stores its neighbor's color in Variable new_v, and updates its R_v to R_u. Yet, it does not change its parent. This behavior creates a path (thanks to Variable new) between a root r of a sub spanning tree $\mathtt{T}_r$ and a node contained in an other sub spanning tree $\mathtt{T}_{r'}$ rooted in r', with $\mathsf{R}_{r'} > \mathsf{R}_r$ (see rule $\mathbb{R}_{\texttt{Path}}$). The subtree $\mathtt{T}_r$ is then re-rooted toward a node aware of a root with a bigger signature u. Now, when $v \in \mathtt{T}_r$'s neighbor u becomes root, it takes u as

a parent (see rules $\mathbb{R}_{\texttt{ReRoot}}$ and $\mathbb{R}_{\texttt{Del}}$). Finally, the descendants of the re-rooted root update their root variables (see rule $\mathbb{R}_{\texttt{Update}}$). The predicate $\mathrm{ErST}(v)$ captures trivial errors and impostor-root

Algorithm 3: Algorithm **ST**

$$
\begin{array}{lll}
\mathbb{R}_{\texttt{Del}} & : (\mathsf{p}_{\mathsf{p}_v} = \mathsf{c}_v) & \longrightarrow \mathsf{p}_v := \varnothing; \\
\mathbb{R}_{\texttt{Update}} & : (\mathsf{p}_{\mathsf{p}_v} \neq \mathsf{c}_v) \wedge (f(v) \neq \varnothing) \wedge (\mathsf{R}_v < \mathsf{R}_{f(v)}) \wedge (\mathsf{p}_v = f(v)) \wedge (\mathsf{new}_v = \varnothing) & \longrightarrow \mathsf{R}_v := \mathsf{R}_f; \\
\mathbb{R}_{\texttt{Path}} & : (\mathsf{p}_{\mathsf{p}_v} \neq \mathsf{c}_v) \wedge (f(v) \neq \varnothing) \wedge (\mathsf{R}_v < \mathsf{R}_{f(v)}) \wedge (\mathsf{p}_v \notin \{\varnothing, f(v)\}) \wedge (\mathsf{new}_v = \varnothing) & \longrightarrow (\mathsf{R}_v, \mathsf{new}_v) := (\mathsf{R}_{f(v)}, f(v)); \\
\mathbb{R}_{\texttt{Merge}} & : (\mathsf{p}_v = \varnothing) \wedge (f(v) \neq \varnothing) \wedge (\mathsf{R}_v < \mathsf{R}_{f(v)}) \wedge (\mathsf{new}_{f(v)} = \varnothing) & \longrightarrow (\mathsf{p}_v, \mathsf{R}_v) := (f(v), \mathsf{R}_f); \\
\mathbb{R}_{\texttt{ReRoot}} & : (\mathsf{p}_v = \varnothing) \wedge (f(v) \neq \varnothing) \wedge (\mathsf{R}_v = \mathsf{R}_{f(v)}) \wedge (\mathsf{new}_v \neq \varnothing) & \longrightarrow (\mathsf{p}_v, \mathsf{new}_v) := (\mathsf{new}_v, \varnothing);
\end{array}
$$

errors for the construction of the spanning tree, these errors are formalized in predicate $\mathrm{ErST}(v)$ formalized in the appendix. Note that predicate $\mathrm{ErST}(v)$ is used in **Freeze** only (and not in **ST**) as these errors are never created by **ST** and **Freeze** has higher priority than **Freeze** (see Section 3.5).

Theorem 5. *Algorithm* **ST** *solves the spanning tree-construction problem in a silent self-stabilizing manner in any n-node graph, assuming the absence of spanning cycle, the state model, and a distributed unfair scheduler, using $O(\log n)$ bits of memory per node.*

3.5.1 Spanning tree-construction with compact identifiers

We adapt our algorithm **ST** to use compact identifiers and obtain Algorithm **C-ST**. It is simple to compare two compact identifiers when the nodes are neighbors. Yet, along the algorithm execution, some nodes become non-root, and therefore the remaining root of fragments can be far away, separated by non-root nodes. To enable multi-hop comparison, we use a broadcasting and convergecast wave on a spanning structure. Let v a node that wants to broadcast its compact identifier. We add an variable check to our previous algorithm. This variable checks whether every descendant or neighbor shares the same compact identifier at the same phase before proceeding to the convergecast. More precisely, a node u must checks if every neighbors w has $\mathcal{C}id_u \simeq_c \mathcal{C}id_w$, and if every child has $\mathsf{check}_v = true$. If so, it sets its variable $\mathsf{check}_v = true$, and the process goes on until node v. As a consequence, v increases or restarts its phase and assigns $false$ to check.

Theorem 6. *Algorithm* **C-ST** *solves the spanning tree-construction problem in a talkative self-stabilizing manner in any n-node graph, assuming the absence of spanning cycle, the state model, and a distributed unfair scheduler, in $O(\max\{\log \Delta, \log\log n\})$ bits of memory per node.*

4 Self-stabilizing leader election

We now present the final assembly of tools we developed to obtain a self-stabilizing leader election algorithm. We add to Algorithm **C-ST** an extra variable ℓ that is mainntained as follows: if a node v has no parent, then $\ell_v = true$, otherwise, $\ell_v = false$. Variable ℓ_v is meant to be the output of the leader election process.

Proof sketch of Theorem 1: Our self-stabilizing leader election algorithm results from combining severals algorithms. As already explained, a higher priority algorithm resets all the variables used by lesser priority algorithms. Moreover, lesser priority algorithm do not modify the variables of the higher priority algorithms. Algorithms are prioritized as follows: **C-Color**, **Freeze**, **C-Break** and **C-ST**. Only the algorithm **C-ST** is talkative, we first proof that the number of activations of rules of algorithm **C-ST** are bounded if there exist nodes enabled by **C-Color**, **Freeze** or **C-Break**. So we already proof the convergence of algorithms **C-Color**, **Freeze** and **C-Break**. Thanks to Theorem 6, we obtain a spanning tree rooted in the node with the maximum degree, maximum color, and maximum identifier. As a consequence, only the root r has $\ell_r = true$ and every other node $v \in V \backslash \{r\}$ has $\ell_v = false$.

References

[1] J. Adamek, M. Nesterenko, and S. Tixeuil. Using abstract simulation for performance evaluation of stabilizing algorithms: The case of propagation of information with feedback. In *SSS 2012*, LNCS. Springer, 2012.

[2] Y. Afek and A. Bremler-Barr. Self-stabilizing unidirectional network algorithms by power supply. *Chicago J. Theor. Comput. Sci.*, 1998.

[3] Yehuda Afek, Shay Kutten, and Moti Yung. Memory-efficient self stabilizing protocols for general networks. In *Distributed Algorithms, 4th International Workshop, WDAG '90, Bari, Italy, September 24-26, 1990, Proceedings*, pages 15–28, 1990.

[4] A. Arora and M. G. Gouda. Distributed reset. *IEEE Trans. Computers*, 43(9):1026–1038, 1994.

[5] B. Awerbuch and R. Ostrovsky. Memory-efficient and self-stabilizing network reset. In *PODC*, pages 254–263. ACM, 1994.

[6] J. Beauquier, M. Gradinariu, and C. Johnen. Memory space requirements for self-stabilizing leader election protocols. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC 1999)*, pages 199–208, 1999.

[7] J. R. S. Blair and F. Manne. An efficient self-stabilizing distance-2 coloring algorithm. *Theor. Comput. Sci.*, 444:28–39, 2012.

[8] L. Blin, F. Boubekeur, and S. Dubois. A self-stabilizing memory efficient algorithm for the minimum diameter spanning tree under an omnipotent daemon. In *2015 IEEE International Parallel and Distributed Processing Symposium, IPDPS*, pages 1065–1074, 2015.

[9] L. Blin, S. Dolev, M. Gradinariu Potop-Butucaru, and S. Rovedakis. Fast self-stabilizing minimum spanning tree construction - using compact nearest common ancestor labeling scheme. In *Distributed Computing, 24th International Symposium, DISC 2010*, pages 480–494, 2010.

[10] L. Blin and P. Fraigniaud. Space-optimal time-efficient silent self-stabilizing constructions of constrained spanning trees. In *35th IEEE International Conference on Distributed Computing Systems, ICDCS 2015*, pages 589–598, 2015.

[11] L. Blin, M. Potop-Butucaru, and S. Rovedakis. A super-stabilizing $\log(n)\log(n)$-approximation algorithm for dynamic steiner trees. *Theor. Comput. Sci.*, 500:90–112, 2013.

[12] L. Blin and S. Tixeuil. Compact deterministic self-stabilizing leader election: The exponential advantage of being talkative. In *Proceedings of the 27th International Conference on Distributed Computing (DISC 2013)*, Lecture Notes in Computer Science (LNCS), pages 76–90. Springer Berlin / Heidelberg, 2013.

[13] L. Blin and S. Tixeuil. Compact deterministic self-stabilizing leader election on a ring: The exponential advantage of being talkative. *Distributed Computing*, page To be appear, 2017.

[14] N.S. Chen, H.P. Yu, and S.T. Huang. A self-stabilizing algorithm for constructing spanning trees. *Information Processing Letters*, 39(3):147 – 151, 1991.

[15] Z. Collin and S. Dolev. Self-stabilizing depth-first search. *Information Processing Letters*, 49(6):297 – 301, 1994.

[16] S. Delaët, B. Ducourthial, and S. Tixeuil. Self-stabilization with r-operators revisited. *Journal of Aerospace Computing, Information, and Communication (JACIC)*, 3(10):498–514, 2006.

[17] E. W. Dijkstra. Self-stabilizing systems in spite of distributed control. *Commun. ACM*, 17(11):643–644, 1974.

[18] S. Dolev. *Self-stabilization*. MIT Press, March 2000.

[19] S. Dolev, M. G. Gouda, and M. Schneider. Memory requirements for silent stabilization. *Acta Inf.*, 36(6):447–462, 1999.

[20] S. Dolev, A. Israeli, and S. Moran. Self-stabilization of dynamic systems assuming only read/write atomicity. *Distributed Computing*, 7(1):3–16, 1993.

[21] S. Dubois and S. Tixeuil. A taxonomy of daemons in self-stabilization. Technical Report 1110.0334, ArXiv eprint, October 2011.

[22] B. Ducourthial and S. Tixeuil. Self-stabilization with path algebra. *Theoretical Computer Science (TCS)*, 293(1):219–236, February 2003.

[23] R. G. Gallager, P. A. Humblet, and P. M. Spira. A distributed algorithm for minimum-weight spanning trees. *ACM Trans. Program. Lang. Syst.*, 5(1):66–77, 1983.

[24] M. G. Gouda, J. Arturo Cobb, and C. Huang. Fault masking in tri-redundant systems. In *SSS*, LNCS, pages 304–313. Springer, 2006.

[25] M. Gradinariu and C. Johnen. Self-stabilizing neighborhood unique naming under unfair scheduler. In Rizos Sakellariou, John Keane, John R. Gurd, and Len Freeman, editors, *Euro-Par 2001: Parallel Processing, 7th International Euro-Par Conference Manchester, UK August 28-31, 2001, Proceedings*, volume 2150 of *Lecture Notes in Computer Science*, pages 458–465. Springer, 2001.

[26] T. Herman and S. V. Pemmaraju. Error-detecting codes and fault-containing self-stabilization. *Inf. Process. Lett.*, 73(1-2):41–46, 2000.

[27] T. Herman and S. Tixeuil. A distributed tdma slot assignment algorithm for wireless sensor networks. In *Proceedings of the First Workshop on Algorithmic Aspects of Wireless Sensor Networks (AlgoSensors'2004)*, number 3121 in Lecture Notes in Computer Science, pages 45–58, Turku, Finland, July 2004. Springer-Verlag.

[28] G. Itkis and L. A. Levin. Fast and lean self-stabilizing asynchronous protocols. In *FOCS*, pages 226–239. IEEE Computer Society, 1994.

[29] G. Itkis, C. Lin, and J. Simon. Deterministic, constant space, self-stabilizing leader election on uniform rings. In *WDAG*, LNCS, pages 288–302. Springer, 1995.

[30] J.Beauquier, M. Gradinariu, C. Johnen, and J. O. Durand-Lose. Token-based self-stabilizing uniform algorithms. *J. Parallel Distrib. Comput.*, 62(5):899–921, 2002.

[31] A. Korman, S. Kutten, and T. Masuzawa. Fast and compact self stabilizing verification, computation, and fault detection of an mst. In *Proceedings of the 30th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing*, PODC '11, 2011.

[32] A. J. Mayer, Y. Ofek, R.l Ostrovsky, and M. Yung. Self-stabilizing symmetry breaking in constant-space (extended abstract). In *STOC*, pages 667–678, 1992.

[33] S. Tixeuil. *Algorithms and Theory of Computation Handbook*, pages 26.1–26.45. CRC Press, Taylor & Francis Group, 2009.

A Appendix

A.1 Compact memory using identifiers: Predicates

Node v can trivially detect an error (predicate $\texttt{ErT}(v)$) whenever its compact identifier does not match its global identifier, or its phase is greater than $\widehat{\mathsf{B}}_v$.

$$\texttt{ErT}(v) \equiv \big[\mathcal{C}id_v \neq (\text{Bit}(1, \text{ID}_v), \mathsf{ph}_v, \text{Bit}(\mathsf{ph}_v, \text{ID}_v))\big] \vee (\mathsf{ph}_v > \widehat{\mathsf{B}}_v) \tag{4}$$

Moreover, in normal operation, the phases of neighboring nodes must be close enough: a node's phase may not be more than 1 ahead or behind any of its neighbors; also a node may not have a neighbor ahead and another behind.

Predicate $\texttt{SErB}(v, S)$ captures these conditions, where $S(v)$ denotes a subset of neighbors of v. The set S should be understood as an input provided by an upper layer algorithm.

$$\texttt{SErB}(v, S) \equiv \Big(\exists u, w \in S(v) : \big(\mathsf{ph}_u > \mathsf{ph}_v + 1\big) \vee \big(\mathsf{ph}_u < \mathsf{ph}_v - 1\big) \vee (|\mathsf{ph}_u - \mathsf{ph}_w| = 2)\Big) \tag{5}$$

In a talkative process, node identifiers are published (though compact identifiers) infinitely often. So, when node v and all its active neighbors have reached the maximum phase (*i.e.* $\mathsf{ph}_v = \widehat{\mathsf{B}}_v$), v goes back to phase 1. Then, if v has $\mathsf{ph}_v = \widehat{\mathsf{B}}_v$ and an active neighbor u has $\mathsf{ph}_u = 1$, it is not an error. But if v has $\mathsf{ph}_v = 1$, one active neighbor u has $\mathsf{ph}_u = \widehat{\mathsf{B}}_v$, and another active neighbor w has $\mathsf{ph}_w > 1$, then an error is detected.

$$\begin{aligned}\texttt{TErB}(v, S) \equiv \Big(\exists u, w \in S(v) : \big[(1 < \mathsf{ph}_v < \widehat{\mathsf{B}}_v) \wedge \big((\mathsf{ph}_u > \mathsf{ph}_v + 1\big) \vee \big(\mathsf{ph}_u < \mathsf{ph}_v - 1)\big)\big] \vee \\ \big[(\mathsf{ph}_v = \widehat{\mathsf{B}}_v) \wedge \big((\mathsf{ph}_u > 1) \vee (\mathsf{ph}_u < \widehat{\mathsf{B}}_v - 1) \vee ((\mathsf{ph}_u = \mathsf{ph}_v - 1) \wedge (\mathsf{ph}_w = 1))\big)\big] \vee \\ \big[(\mathsf{ph}_v = 1) \wedge \big((\mathsf{ph}_u > 2) \vee (\mathsf{ph}_u < \widehat{\mathsf{B}}_v) \vee ((\mathsf{ph}_u = \widehat{\mathsf{B}}_v) \wedge (\mathsf{ph}_w = 2))\big)\big]\Big)\end{aligned} \tag{6}$$

If v detects an error through $\texttt{ErT}(v)$, $\texttt{SErB}(v, S)$ or $\texttt{TErB}(v, S)$, it resets its compact identifier to its first phase value:

$$\mathcal{R}esetCid(v) : \mathcal{C}id_v := (\widehat{\mathsf{B}}_v, \mathsf{ph}_v, \mathsf{Bp}_v) = (\text{Bit}(1, \text{ID}_v), 1, \text{Bit}(1, \text{ID}_v)) \tag{7}$$

This may trigger similar actions at neighbors in S, so that all such errors eventually disappear.

The compact identifier of u is smaller (respectively greater) than the compact identifier of v, if the most significant bit-position of u is smaller (respectively greater) than the most significant bit-position of v, or if the most significant bit-position of u is equal to the most significant bit-position of v, u and v are in the same phase, and the bit-position of u is smaller (respectively greater) than the bit-position of v:

$$\mathcal{C}id_u^i <_c \mathcal{C}id_v^i \equiv (\widehat{\mathsf{B}}_u < \widehat{\mathsf{B}}_v) \vee \big((\widehat{\mathsf{B}}_v = \widehat{\mathsf{B}}_u) \wedge (\mathsf{Bp}_u < \mathsf{Bp}_v)\big) \tag{8}$$

When two nodes u and v have the same most significant bit-position and the same bit position at phase $i < \widehat{\mathsf{B}}_v$, they are possibly equal with respect to compact identifiers (denoted by $\simeq_c$).

$$\mathcal{C}id_u^i \simeq_c \mathcal{C}id_v^i \equiv (i < \widehat{\mathsf{B}}_v) \wedge \big((\widehat{\mathsf{B}}_v = \widehat{\mathsf{B}}_u) \wedge (\mathsf{Bp}_u = \mathsf{Bp}_v)\big) \tag{9}$$

Finally, two nodes u and v have the same compact identifier (denoted by $=_c$) if their phase reaches the size of the binary representation of the identifier of the two nodes, and their last bit-position is the same.

$$\mathcal{C}id_u^i =_c \mathcal{C}id_v^i \equiv (i = \widehat{\mathsf{B}}_v = \widehat{\mathsf{B}}_u) \wedge (\mathsf{Bp}_u = \mathsf{Bp}_v) \tag{10}$$

Predicate $\texttt{SPh}^+(v)$ is true if for every node u in $S(v)$, either $\mathcal{C}id_u^i \simeq_c \mathcal{C}id_v^i$, or $\mathsf{ph}_u = \mathsf{ph}_v + 1$.

$$\texttt{SPh}^+(v,S) \equiv \forall u \in S(v) : (\mathcal{C}id_u^i \simeq_c \mathcal{C}id_v^i) \vee (\mathsf{ph}_u = \mathsf{ph}_v + 1) \tag{11}$$

Similarly, $\texttt{TPh}^+(v,S)$ is true if for every node u in $S(v)$, either $\mathcal{C}id_u^i =_c \mathcal{C}id_v^i$, or $\mathsf{ph}_u = 1$.

$$\texttt{TPh}^+(v,S) \equiv \texttt{SPh}^+(v,S) \vee \forall u \in S(v) : (\mathsf{ph}_v = \widehat{\mathsf{B}}_v) \wedge \big((\mathcal{C}id_u^i =_c \mathcal{C}id_v^i) \vee (\mathsf{ph}_u = 1)\big) \tag{12}$$

When $\texttt{TPh}^+(v,S)$ or $\texttt{SPh}^+(v)$ is true, v may increase its phase:

$$\mathcal{I}ncPh(v) : \mathcal{C}id_v := \begin{cases} (\widehat{\mathsf{B}}_v, \mathsf{ph}_v + 1, \mathrm{Bit}(\mathsf{ph}_v + 1, \mathrm{ID}_v)) & \text{if } \mathsf{ph}_v < \widehat{\mathsf{B}}_v \\ (\widehat{\mathsf{B}}_v, 1, \mathrm{Bit}(1, \mathrm{ID}_v)) & \text{if } \mathsf{ph}_v = \widehat{\mathsf{B}}_v \end{cases} \tag{13}$$

In some case, we need to compute the minimum or the maximum on compact identifiers. Let f denote a function that is either minimum or maximum. Let us denote by $\texttt{C}\widehat{\mathsf{B}}(v,S,f)$ the minimum or the maximum most significant bit of nodes in $S(v)$.

$$\texttt{C}\widehat{\mathsf{B}}(v,S,f) = f\{\widehat{\mathsf{B}}_w : w \in S(v)\} \tag{14}$$

To compare compact identifiers, one must always refer to the same phase; we always consider the minimum phase for nodes in $S(v)$.

$$\texttt{CPh}(v,S,f) = \min\{\mathsf{ph}_w : w \in S(v), \widehat{\mathsf{B}}_w = \texttt{C}\widehat{\mathsf{B}}(w,S,f)\} \tag{15}$$

Finally, we compute the minimum or the maximum bit position.

$$\texttt{CBp}(v,S,f) = f\{\mathsf{Bp}_w : w \in S(v), \mathsf{ph}_w = \texttt{CPh}(w,S,f)\} \tag{16}$$

Predicate $\texttt{MinCid}(v,S)$ checks if $\mathcal{C}id_v$ is equal to the minimum among nodes in $S(v)$:

$$\texttt{MinCid}(v,S) \equiv \big(\mathcal{C}id_v = (\texttt{C}\widehat{\mathsf{B}}(v,S,\min), \texttt{CPh}(v,S,\min), \texttt{CBp}(v,S,\min))\big) \tag{17}$$

The predicate $\texttt{MaxCid}(v,S)$ does the same for the maximum:

$$\texttt{MaxCid}(v,S) \equiv \big(\mathcal{C}id_v = (\texttt{C}\widehat{\mathsf{B}}(v,S,\max), \texttt{CPh}(v,S,\max), \texttt{CBp}(v,S,\max))\big) \tag{18}$$

Node v may use $\mathcal{O}pt$ to assign its local variables the minimum (or maximum) compact identifier in $S(v)$.

$$\mathcal{O}pt(v,S,f) : \begin{cases} \widehat{\mathsf{B}}_v := \texttt{C}\widehat{\mathsf{B}}(v,S,f) \\ \mathsf{ph}_v := \texttt{CPh}(v,S,f) \\ \mathsf{Bp}_v := \texttt{CBp}(v,S,f) \end{cases} \tag{19}$$

A.2 Silent Self-stabilizing Distance-2 Coloring: Predicates and Correctness

A.2.1 Predicates

Note that all rules are exclusive, because a node v cannot be both $\texttt{Player}(v)$ and $\texttt{Relay}(v)$. Let us now describe the functions, predicates and actions use by algorithm **C-Color**. Remember that $N[v] = N(v) \cup \{v\}$. Function $\Delta(v)$ returns the maximum degree between v and its neighbors, and is used to define the range $[1, \Delta(v)^2 + 1]$ of authorized colors for a node v:

$$\Delta(v) = \max\{\delta_u : u \in N[v]\} \tag{20}$$

The function $\texttt{mC}(v)$ returns the minimum color in conflict at distance one and two :

$$\texttt{mC}(v) := \min\{\texttt{c}_u : u, w \in N[v] \wedge (u \neq w) \wedge (\texttt{c}_u = \texttt{c}_w)\} \tag{21}$$

The function $\texttt{MC}(v)$ returns the maximum color used at distance one :

$$\texttt{MC}(v) := \max\left\{\texttt{c}_u : u \in N[v]\right\} \tag{22}$$

The predicate $\texttt{Bad}(v)$ is true if v has not yet set the right value for either $\texttt{mC}(v)$ or $\texttt{MC}(v)$. Moreover, this predicate checks if v's compact identifier matches its global identifier (see predicate 4: $\texttt{ErT}(v)$) and if the phases of the subset of v's neighbors $\texttt{Oth}(v)$ are coherent with v (see predicate 5: $\texttt{SErB}(v, \texttt{Oth}(v))$.

$$\texttt{Bad}(v) \equiv (\breve{c}_v \neq \texttt{mC}(v)) \vee (\widehat{c}_v \neq \texttt{MC}(v)) \vee (\texttt{ErT}(v) \vee \texttt{SErB}(v, \texttt{Oth}(v)) \tag{23}$$

The predicate $\texttt{Player}(v)$ is true if v has the minimum color in conflict (announced by its neighbors or by itself). Observe that the conflict may be at distance one or two:

$$\texttt{Player}(v) \equiv (\texttt{c}_v = \min\{\breve{c}_u, u \in N[v]\}) \tag{24}$$

The predicate $\texttt{Relay}(v)$ is true if v does not have the minimum color in conflict, and at least two of its neighbors have the minimum color in conflict:

$$\texttt{Relay}(v) \equiv (\breve{c}_v \neq \texttt{c}_v) \wedge (\breve{c}_v = \min\{\breve{c}_u, u \in N[v]\}) \wedge (\exists u, w \in N(v), (\breve{c}_v = \texttt{c}_u) \wedge (\breve{c}_v = \texttt{c}_w)) \tag{25}$$

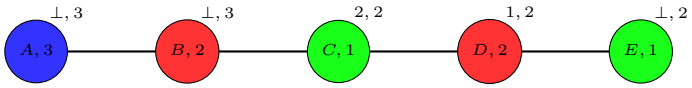

Fig. 2 – The pair in the node are the identifier of the node and the color of the node. The pair outside the node are the minimum color in conflict (or $\perp$ if none) and the maximum color used. Node D is a relay for node C and E because $\breve{c}_D = 1$, the color of C and V. Similarly, Node C is a relay for nodes B and D.

The function $\texttt{PlayR}(v)$ returns the subset of v's neighbors that have the minimum color in conflict, when v is a relay node:

$$\texttt{PlayR}(v) := \{u : u \in N(v) \wedge \texttt{c}_u = \breve{c}_v\} \tag{26}$$

The function $\texttt{Oth}(v)$ returns the subset of v's neighbors that are in conflict with v at distance one, or the set of relay nodes for the conflict at distance two, when v has $\texttt{Player}(v)$ equal to *true*:

$$\texttt{Oth}(v) := \{u : u \in N(v) \wedge \texttt{c}_u = \texttt{c}_v\} \cup \{u : u \in N(v) \wedge \breve{c}_u = \texttt{c}_v\} \tag{27}$$

The predicate $\texttt{Loser}(v)$ is true whenever a competing player of v has a greater bit position at the same phase. A node whose identifier is maximum among competitors does not change its color, but any loosing competitor does.

$$\texttt{Loser}(v) \equiv \exists u \in \texttt{Oth}(v) : \mathcal{C}id_v^i <_c \mathcal{C}id_u^i \tag{28}$$

The predicate $\texttt{RUp}(v)$ is true if a relay node is not according to its player neighbors, like we decide to change the color of the node with the minimum identifier the relay node stores the maximum compact identifier of its player neighbors:

$$\texttt{RUp}(v, \texttt{PlayR}(v)) \equiv \mathcal{C}id_v \neq_c \texttt{MaxCid}(v, \texttt{PlayR}(v)) \tag{29}$$

The action $\mathcal{U}pdate(v)$ updates the variables $\check{c}_v, \hat{c}_v$ and resets the variables relatives to the identifier (see command $\mathcal{R}esetCid$(v) in equation 7).

$$\mathcal{U}pdate(v) : \check{c}_v := \texttt{mC}(v); \hat{c}_v := \texttt{MC}(v); \mathcal{R}esetCid(v); \tag{30}$$

When a node change its color, it takes the maximum color at distance one and two plus one modulo $\Delta(v)^2 + 1$, and then add one to assign colors in the range $[1, \ldots, \Delta(v)^2 + 1]$.

$$\mathcal{N}ewcolor(v) : \mathsf{c}_v := \big((\max\{\texttt{MC}_u : u \in N[v]\} + 1) \mod \Delta(v)^2 + 1\big) + 1; \tag{31}$$

A.2.2 Correctness

In the details of lemmas that are presented in the sequel, we use predicates on configurations. These predicates are mean to be intermediate attractors towards a legitimate configuration (*i.e.*, a configuration with a unique leader). To establish that those predicates are indeed attractors, we use potential functions [33], that is, functions that map configurations to non-negative integers, and that strictly decrease after any algorithm step is executed.

To avoid additional notations, we use sets of configurations to define predicates; the predicate should then be understood as the characteristic function of the set (that returns *true* if the configuration is in the set, and *false* otherwise).

Lemma 1. *Using a range of* $[1, \Delta(v)^2 + 1]$ *for colors at node* v *is sufficient to enable distance-2 coloring of the graph.*

Proof. In the worst case, all neighbors at distance one and two of v have different colors. Now, v has at most $\Delta(v)$ neighbors at distance one, each having $\Delta(v) - 1$ other neighbors than v. In total, v has at most $\Delta(v)^2 - \Delta(v)$ neighbors at distance up to two, each having a distinct color. Using a range of $[1, \Delta(v)^2 + 1]$ for v's color leaves at least $\Delta(v) + 1$ available colors for node v. □

Let $\lambda : \Gamma \times V \to \mathbb{N}$ be the following function:

$$\lambda(\gamma, v) = \begin{cases} 2 & \text{if } \delta_v \neq \deg(v) \text{in} \gamma \\ 1 & \text{if } (\delta_v = \deg(v) \text{in} \gamma) \wedge \mathsf{c}_v(\gamma) > \Delta(v)^2 + 1 \\ 0 & \text{otherwise} \end{cases}$$

Let $\Lambda : \Gamma \to \mathbb{N}$ be the following potential function:

$$\Lambda(\gamma) = \sum_{v \in V} \lambda(\gamma, v)$$

Let $\tau : \Gamma \times V \to \mathbb{N}$ be the following function:

$$\tau(\gamma, v) = \begin{cases} \Delta(v)^2 + 1 + \mathsf{c}_v(\gamma_0) - \mathsf{c}_v(\gamma) & \text{if } \mathsf{c}_v(\gamma) > \mathsf{c}_v(\gamma_0) \\ \mathsf{c}_v(\gamma_0) - \mathsf{c}_v(\gamma) & \text{otherwise} \end{cases}$$

Where $\mathsf{c}_v(\gamma_0)$ is the color of v in configuration γ_0, γ_0 being defined as the configuration where $\Lambda(\gamma_0, v)$ reaches zero. Also, $\mathsf{c}_v(\gamma)$ is the color of v in configuration γ.

Let $\mathcal{C} : \Gamma \to \mathbb{N}$ be the following potential function:

$$\mathcal{C}(\gamma) = (\kappa_0(\gamma), \kappa_1(\gamma), \ldots, \kappa_{\Delta^2}(\gamma))$$

where $\kappa_i(\gamma) = |\{v \in V : \tau(\gamma, v) = i\}|$. The comparison between two configurations $\mathcal{C}(\gamma)$ and $\mathcal{C}(\gamma')$ is done using lexical ordering. We denote by γ' the configuration after activation of (a subset of) the nodes in $A_\kappa(\gamma)$ where $A_\kappa(\gamma)$ denotes the enabled nodes in γ due to rule $\mathbb{R}_{\texttt{Color}}$. We can now prove the following result: $\mathcal{C}(\gamma') < \mathcal{C}(\gamma)$ for every configuration γ where $A_\kappa(\gamma)$ is not empty.

Lemma 2. *$\mathcal{C}(\gamma') < \mathcal{C}(\gamma)$ for every configuration γ such that $A_\kappa(\gamma)$ is not empty.*

Proof. We consider a node $v \in A_\kappa(\gamma)$. After executing rule $\mathbb{R}_{\texttt{Color}}$, v takes a color $\mathsf{c}_v(\gamma') = \big((\max\{\texttt{MC}_u : u \in N[v]\} + 1) \mod \Delta(v)^2 + 1\big) + 1$. As a consequence $\tau(\gamma, v)$ decreases by one, so $\mathcal{C}(\gamma') < \mathcal{C}(\gamma)$. □

Let $\psi : \Gamma \times V \to \mathbb{N}$ be the function defined by:

$$\psi(\gamma, v) = \begin{cases} n & \text{if } \texttt{Bad}(v) \vee (\widehat{\mathsf{B}}_v, \mathsf{ph}_v, \mathsf{Bp}_v) \text{ is true} \\ 2\log n - \mathsf{ph}_v & \text{if } \texttt{Player}(v) \text{ is true} \\ 2 & \text{if } \texttt{Relay}(v) \wedge \neg\texttt{RUp}(v, \texttt{PlayR}(v)) \text{ is true} \\ 1 & \text{if } \texttt{Relay}(v) \wedge \texttt{RUp}(v, \texttt{PlayR}(v)) \text{ is true} \\ 0 & \text{otherwise} \end{cases}$$

Let $\Psi : \Gamma \to \mathbb{N}$ be the potential function defined by:

$$\Psi(\gamma) = \sum_{v \in V} \psi(\gamma, v)$$

Let $\Phi : \Gamma \to \mathbb{N}$ be the potential function defined by:

$$\Phi(\gamma) = (\Lambda(\gamma), \mathcal{C}(\gamma), \Psi(\gamma)).$$

The comparison between two configurations $\Phi(\gamma)$ and $\Phi(\gamma')$ is by using lexical order. We denote by $A(\gamma)$ the (subset of) enabled nodes (for any rule of our algorithm) in configuration γ. Note that the algorithm is stabilized when every node is neither a player nor a relay, that is the nodes have no conflict at distance one and two, when $\Psi(\gamma) = 0$. We define

$$\Gamma_{\mathcal{C}} = \{\gamma \in \Gamma : \Phi(\gamma) = 0\}$$

Lemma 3. *$true \rhd \Gamma_{\mathcal{C}}$ and $\Gamma_{\mathcal{C}}$ is closed.*

Proof. The function $\Lambda(\gamma)$ decreases by any execution of rules $\mathbb{R}_\Delta$ and $\mathbb{R}^+_\Delta$. Remark that $\deg(v)$ is considered a non corruptible local information, so once v has executed $\mathbb{R}_\Delta$, this rule remains disabled afterwards. Moreover, $\mathbb{R}_{\texttt{Color}}$ maintains the value of the color inferior (or equal) to $\Delta(v)^2 + 1$, and other rules modifying the color maintain this invariant. Hence, if the scheduler activates rules $\mathbb{R}_\Delta$ or $\mathbb{R}^+_\Delta$, we obtain $\Lambda(\gamma') < \Lambda(\gamma)$, otherwise if other rules are activated, then $\Lambda(\gamma') = \Lambda(\gamma)$. We already saw that, when the scheduler activates a node v for rule $\mathbb{R}_{\texttt{Color}}$, we obtain $\mathcal{C}(\gamma') < \mathcal{C}(\gamma)$. Overall, if the scheduler activates rules $\mathbb{R}_\Delta$, $\mathbb{R}^+_\Delta$, or $\mathbb{R}_{\texttt{Color}}$ we obtain $\Phi(\gamma') < \Phi(\gamma)$. We now consider the cases where the scheduler activates other rules.

First, we focus on rule $\mathbb{R}_{\texttt{Up}}$. Let us consider A', the set of nodes enabled for this rule, and a node v such that $v \in A'$. Then, v has $\texttt{Bad}(v) = \mathit{true}$ (see predicate 23), or $(\texttt{Player}(v) \wedge (\texttt{ErT}(v) \vee \texttt{SErB}(v, \texttt{Oth}(v))))$ is *true*. If v has $(\check{c}_v \neq \texttt{mC}(v)) \vee (\hat{c}_v \neq \texttt{MC}(v))$ in γ, then after activation of v, we obtain $\check{c}_v = \texttt{mC}(v)$ and $\hat{c}_v = \texttt{MC}(v)$ in γ' because $\texttt{mC}(v)$ and $\texttt{MC}(v)$ depend only on the color of the neighbors of v (see Function 21 and 22). The same argument applies for $(\texttt{ErT}(v) \vee \texttt{SErB}(v, \texttt{Oth}(v)))$, because $\mathcal{C}id_v$ is computed only with the identifier of v. So, after execution of $\mathbb{R}_{\texttt{Up}}$ by v, we obtain $\psi(\gamma', v) < \psi(\gamma, v)$. Remark that, if the color of the neighbors of v does not change, rule $\mathbb{R}_{\texttt{Up}}$ remains disabled. Now, if the color changes, $\Phi(\gamma)$ still decreases thanks to Lemma 2.

Let us consider now a configuration where the rule $\mathbb{R}_{\texttt{Up}}$ is disabled for every node. Rule $\mathbb{R}_{\texttt{Bit}}$ increases the phase of a player node, so after activation of this rule we obtain $\psi(\gamma', v) < \psi(\gamma, v)$. Executing rule $\mathbb{R}_{\texttt{Relay}}$ decreases also $\psi(\gamma', v)$ due to $\texttt{RUp}(v, \texttt{PlayR}(v))$, because when all nodes in $\texttt{PlayR}(v)$ have increase their phases, ψ decreases for all v's neighbors. □

Lemma 4. *Algorithm* **C-Color** *requires* $O(\max\{\log \Delta, \log\log n\})$ *bits of memory per node.*

Proof. The variables $\delta_v, \mathsf{c}_v, \check{c}_v, \hat{c}_v$ take $O(\log \Delta)$ bits. The compact identifier $\mathcal{C}id_v$ takes $O(\log\log n)$ bits per node. □

Proof of Theorem 2. Direct by Lemma 1, Lemma 3 and Lemma 4. □

Lemma 5. *Algorithm* **C-Color** *converges in* $O(\Delta^{\Delta^2} n^3)$ *steps.*

Proof. Direct by the potential function $\Phi(\gamma)$. □

A.3 Cleaning a cycle or an impostor-rooted spanning tree

We now present **Freeze** in Algorithm 4. This algorithm uses only one binary variable froz. This approach presents several advantages. After v detecting a cycle, the cycle is broken (v deletes its parent), and a frozen node cannot reach its own subtree, due to the cleaning process taking place from the leaves to the root. So, two cleaning processes cannot create a livelock.

Algorithm 4: Algorithm **Freeze**

$$
\begin{array}{lcll}
\mathbb{R}_{\texttt{Error}} & : & \mathrm{ErCycle}(v) \vee \mathrm{ErST}(v) & \longrightarrow \mathsf{froz}_v := 1, \mathsf{p}_v := \varnothing; \\
\mathbb{R}_{\texttt{Froze}} & : & \neg\mathrm{ErCycle}(v) \wedge \neg\mathrm{ErST}(v) \wedge (\mathsf{froz}_{\mathsf{p}_v} = 1) \wedge (\mathsf{froz}_v = 0) & \longrightarrow \mathsf{froz}_v := 1; \\
\mathbb{R}_{\texttt{Prun}} & : & \neg\mathrm{ErCycle}(v) \wedge \neg\mathrm{ErST}(v) \wedge (\mathsf{froz}_{\mathsf{p}_v} = 1) \wedge (\mathsf{froz}_v = 1) \wedge (\texttt{Ch}(v) = \varnothing) & \longrightarrow \mathit{Reset}(v);
\end{array}
$$

Theorem 7. *Algorithm* **Freeze** *deletes a cycle or an impostor-rooted sub spanning tree in n-nodes graph in a silent self-stabilizing manner, assuming the state model, and a distributed unfair scheduler. Moreover, Algorithm* **Freeze** *uses* $O(1)$ *bits of memory per node.*

Lemma 6. *Algorithm* **Freeze** *converges in* $O(n)$ *steps.*

Proofs of Theorem 7 and Lemma 6 see article [13].

A.4 Self-stabilizing Cycle Detection

A.4.1 Correctness of the algorithm Break

Let I_i be the nearest descendant of v ($I_i \neq v$) such that $\mathsf{m}_{I_i} = i$, if such a node exists. And let us denote by $D(v,i)$ the set of nodes on the path between v and I_i. We suppose that every node u in $D(v,i)$ has ErCycle(u) = $false$. The value i can improve the value m_v if and only every node u in $D(v,i)$ has a $\mathsf{m}_u > i$ and $\mathsf{E}_u \neq i$. Also, if $\mathsf{E}_u = i$, the value vanishes during the execution, note that may be $u = v$. Predicate Improve captures this fact.

$$\begin{aligned}\text{Improve}(v,i) \equiv \big(\forall u \in D(v,i) : (\mathsf{m}_u > i)\big) \wedge \big[\big(\forall u \in D(v,i) : \mathsf{E}_u \neq i)\big) \vee \\ \big(\exists u \in D(v,i) : (\mathsf{E}_u = i) \wedge \text{Improve}(u,j)\big)\big]\end{aligned} \tag{32}$$

Let $\alpha : \Gamma \times V \to \mathbb{N}$ be the function defined by:

$$\alpha(\gamma,v,i) = \begin{cases} 4 & \text{if } (\mathsf{m}_v > i) \wedge (\mathsf{E}_v \neq i) \wedge \text{Improve}(v,i) \\ 3 & \text{if } (\mathsf{m}_v = i) \wedge (\mathsf{E}_v \neq i) \wedge \text{Improve}(v,i) \\ 2 & \text{if } (\mathsf{m}_v = i) \wedge (\mathsf{E}_v = i) \wedge \text{Improve}(v,i) \\ 1 & \text{if } (\mathsf{m}_v > i) \wedge (\mathsf{E}_v = i) \wedge \text{Improve}(v,i) \\ 0 & \text{if } \neg\text{Improve}(v,i) \end{cases}$$

Let $\beta : \Gamma \times V \to \mathbb{N}$ be the function defined by:

$$\beta(\gamma,i) = \sum_{v \in V} \alpha(\gamma,v,i)$$

Let $\Xi : \Gamma \to \mathbb{N}$ be the function defined by:

$$\Xi(\gamma) = (\beta(\gamma,0), \beta(\gamma,1), \ldots, \beta(\gamma,IdMax))$$

The comparaison between two configurations $\Xi(\gamma)$ and $\Xi(\gamma')$ is performed using lexical order. In the following, $\mathsf{m}_v(\gamma)$ denotes the variable m_v in configuration γ. Note that the algorithm is stabilized when no value i can improve the value stored in m_v, that is when $\Xi(\gamma) = 0$. We define

$$\Gamma_{\mathcal{B}} = \{\gamma \in \Gamma : \Xi(\gamma) = 0\}$$

Lemma 7. *$true \rhd \Gamma_{\mathcal{B}}$ and $\Gamma_{\mathcal{B}}$ is closed.*

Proof.
- Rule $\mathbb{R}_{\texttt{Start}}(v)$: $\mathsf{m}_v(\gamma) = \mathsf{m}_v(\gamma')$, so for $i < \mathsf{m}_v(\gamma')$, we have $\beta(\gamma',i) = \beta(\gamma,i)$. Note that for $i > \mathsf{m}_v(\gamma)$, v has no effect on $\beta(\gamma,i)$ and $\beta(\gamma',i)$. Now, $\beta(\gamma,\mathsf{m}_v) = 3$ because $\mathbb{R}_{\texttt{Start}}(v)$ is enabled for v only if $(\mathsf{E}_v \neq \mathsf{m}_v)$, and $\beta(\gamma',\mathsf{m}_v) = 2$ because we have $(\mathsf{E}_v = \mathsf{m}_v)$, thus $\beta(\gamma,\mathsf{m}_v) = 3 > \beta(\gamma',\mathsf{m}_v) = 2$. So, if the scheduler activates v with rule $\mathbb{R}_{\texttt{Start}}(v)$, we obtain $\Xi(\gamma') < \Xi(\gamma)$.
- Rule $\mathbb{R}_{\texttt{ID}}$:
  - $i < \mathsf{m}_v(\gamma)$: $\beta(\gamma',i) = \beta(\gamma,i)$ because $\text{ID}_v > \mathsf{m}_v(\gamma)$ (otherwise an error is detected). Also, if i can improve $\mathsf{m}_v(\gamma)$, it can also improve $\mathsf{m}_v(\gamma')$.
  - $\mathsf{m}_v(\gamma)$**:** Rule $\mathbb{R}_{\texttt{ID}}$ needs $\mathsf{E}_v(\gamma) = \mathsf{m}_v(\gamma)$, so $\beta(\gamma,\mathsf{m}_v(\gamma)) = 2$. Now, we have $\mathsf{E}_v(\gamma') = \mathsf{m}_v(\gamma) \neq \mathsf{m}_v(\gamma')$, and we obtain $\beta(\gamma,\mathsf{m}_v(\gamma)) = 2 > \beta(\gamma',\mathsf{m}_v(\gamma)) = 1$.

  As a consequence, $\beta(\gamma',i) = \beta(\gamma,i)$ for $i < \mathsf{m}_v(\gamma)$ and $\beta(\gamma',\mathsf{m}_v(\gamma)) < \beta(\gamma,\mathsf{m}_v(\gamma))$. So, if the scheduler activates only v for rule $\mathbb{R}_{\texttt{ID}}(v)$, we obtain $\Xi(\gamma') < \Xi(\gamma)$.

- Rule $\mathbb{R}_{\texttt{Min}}(v)$: $\mathsf{m}_v(\gamma') < \mathsf{m}_v(\gamma)$ and $\mathsf{E}_v(\gamma) \neq \mathsf{m}_v(\gamma')$ by definition of rule $\mathbb{R}_{\texttt{Min}}(v)$, so in γ, we have $\alpha(\gamma, v, \mathsf{m}_v(\gamma')) = 4$ because $\mathsf{m}_v(\gamma') < \mathsf{m}_v(\gamma)$ and $\mathsf{m}_v(\gamma')$ can improve $\mathsf{m}_v(\gamma)$. Now, we have $\alpha(\gamma', v, \mathsf{m}_v(\gamma')) = 3$ because $\mathsf{E}_v(\gamma') \neq \mathsf{m}_v(\gamma')$. $\mathsf{m}_v(\gamma') < \mathsf{m}_v(\gamma)$, so if the scheduler activates only v for rule $\mathbb{R}_{\texttt{Min}}(v)$, we obtain $\Xi(\gamma') < \Xi(\gamma)$.

To conclude, $\Xi(\gamma') < \Xi(\gamma)$ for every configurations γ and γ', when γ' occurs later than γ. □

Proof of Theorem 3. Now we prove that, in $\Gamma_{\mathcal{B}}$ if the spanning structure S contains a least one cycle $\mathtt{C}$, then at least one node v in $\mathtt{C}$ has ErCycle$(v) = true$. For the purpose of contradiction, let us assume the opposite. Let $\gamma \in \Gamma_{\mathcal{B}}$ and every node v in the cycle $\mathtt{C}$ in γ has ErCycle$(v) = false$. By definition all the nodes in $\mathtt{C}$ have a parent, and all the nodes have $\mathsf{m}_v \geqslant \mathsf{m}_{\mathsf{k}_v}$. Now, if a node v shares the same m with its parent and its child, then rule $\mathbb{R}_{\texttt{Start}}$ is enabled for v, a contradiction with $\Xi(\gamma) = 0$. In a cycle, it is not possible that all nodes have $\mathsf{m}_v > \mathsf{m}_{\mathsf{k}_v}$ (due to well foundedness of integers, at least one node v has $\mathsf{m}_v < \mathsf{m}_{\mathsf{k}_v}$), which is a contradiction with the assumption that every node v is such that ErCycle$(v) = false$. □

Lemma 8. *Algorithm* **Break** *converges in* $O(n^n)$ *steps.*

Proof. Direct by the potential function $\Xi(\gamma)$. □

A.4.2 Algorithm C-Break et Predicates

Predicate ErCycle must be updated to take into account this extra care. We denote by k_v the child of v with minimum compact identifier stored in m_v. Moreover, predicate ErCycle now takes into account the error(s) related to compact identifiers management (see Equation 4: $\texttt{ErT}(v)$). It is important to note that the cycle breaking algorithm does not manage phase differences. Indeed, a node v with a phase bigger than the phase of one of its children takes the m of its children, if its phase its bigger than two or if no child has its same compact identifier. The compact identifier stored in m_v is be compared using lexicographic order by rule $\mathbb{R}_{\texttt{Min}}$.

$$\begin{aligned}\text{ErCycle}(v) \equiv (\mathsf{p}_v \neq \varnothing) \wedge \Big((\mathsf{m}_{\mathsf{k}(v)} =_c \mathcal{C}id_v^i) \vee (\exists (u,w) \in \texttt{Ch}(v) : \mathsf{m}_u =_c \mathsf{m}_w) \vee (\mathsf{m}_v >_c \mathcal{C}id_v^i) \vee \\ \big((\mathsf{m}_v \not\approx_c \mathcal{C}id_v) \wedge (\mathsf{m}_v <_c \mathsf{m}_{\mathsf{k}(v)})\big) \vee (\texttt{Active}_v \wedge \texttt{ErT}(v))\Big)\end{aligned} \tag{33}$$

Algorithm 5: Algorithm **C-Break** For node v with $\neg$ErCycle(v)

$$\begin{array}{lll}
\mathbb{R}_{\texttt{Inc}} & : \texttt{Active}_v \wedge (\mathsf{m}_{\mathsf{p}_v} \simeq_c \mathsf{m}_v \simeq_c \mathsf{m}_{\mathsf{k}(v)}) & \longrightarrow \mathcal{I}ncPh(v); \\
\mathbb{R}_{\texttt{Start}} & : \neg\texttt{Active}_v \wedge (\mathsf{m}_{\mathsf{p}_v} \simeq_c \mathsf{m}_v \simeq_c \mathsf{m}_{\mathsf{k}(v)}) \wedge (\mathsf{E}_v \not\approx_c \mathsf{m}_v) & \longrightarrow \mathsf{E}_v := \mathsf{m}_v; \\
\mathbb{R}_{\texttt{Min}} & : (\mathsf{m}_v >_c \mathsf{m}_{\mathsf{k}(v)}) \wedge (\mathsf{E}_v \not\approx_c \mathsf{m}_{\mathsf{k}(v)}) & \longrightarrow \mathsf{m}_v := \mathsf{m}_{\mathsf{k}(v)}, \texttt{Active}_v := false; \\
\mathbb{R}_{\texttt{ID}} & : (\mathsf{E}_{\mathsf{p}_v} \simeq_c \mathsf{E}_v \simeq_c \mathsf{E}_{\mathsf{k}(v)} \simeq_c \mathsf{m}_v) \wedge (\mathsf{m}_v \not\approx_c \mathcal{C}id_v^1) & \longrightarrow \mathsf{m}_v := \mathcal{C}id_v^1, \texttt{Active}_v := true;
\end{array}$$

The proof of theorem 4 mimics the proof of algorithm **Break**.

Lemma 9. *Algorithm* **C-Break** *converges in* $O(n^n \log n)$ *steps.*

A.5 Spanning Tree Construction without Distance to the Root Maintenance

A.5.1 Predicates

The function $f(v)$ returns the color of the neighbor of v with the maximum root:

$$f(v) = \{\mathsf{c}_u : u \in N(v) \wedge \mathsf{R}_u = \max\{\mathsf{R}_w : w \in N(v)\}\} \tag{34}$$

We now present a list of trivial errors and impostor-root errors for the construction of the spanning tree. The explanations of the different elements composing the predicate $\mathrm{ErST}(v)$ follow: (1) A node without relative has its root signature different to its own variables. (2) The variable δ_v is not equal at the degree of v. (3) The invariant is not satisfied. (4) A node with $\mathsf{new}_v = \varnothing$ (that is, v is not involved in a rerouting process) has a root signature bigger than that of its parent;(5) A node with $\mathsf{new}_v \neq \varnothing$ (that is, v is involved in a rerouting process) has a root signature different from that of its tentative new parent.(6) A root v with $\mathsf{new}_v = \varnothing$ and a signature R_v that does not match is own.(7)A node involved in a rerouting process whose parent's parent is itself.

$$\mathrm{ErST}(v) \equiv \begin{cases} (1) & \big((\mathsf{p}_v = \varnothing) \wedge (\mathtt{Ch}(v) = \varnothing) \wedge [(\mathsf{R}_v \neq (\deg_v, \mathsf{c}_v, \mathrm{ID}_v)) \vee (\mathsf{new}_v \neq \varnothing)]\big) \vee \\ (2) & (\delta_v \neq \deg_v) \vee \\ (3) & (\overline{\mathsf{n}}(v) \notin \{\{\mathsf{p}_v\} \cup \mathtt{Ch}(v)\}\}) \vee \\ (4) & \big((\mathsf{new}_v = \varnothing) \wedge (\mathsf{R}_v > \mathsf{R}_{\mathsf{p}_v})\big) \vee \\ (5) & \big((\mathsf{new}_v \neq \varnothing) \wedge (\mathsf{R}_v \neq \mathsf{R}_{\mathsf{new}_v})\big) \vee \\ (6) & \big((\mathsf{p}_v = \varnothing) \wedge (\mathsf{new}_v = \varnothing) \wedge (\mathsf{R}_v \neq (\deg_v, \mathsf{c}_v, \mathrm{ID}_v))\big) \vee \\ (7) & \big((\mathsf{new}_v \neq \varnothing) \wedge (\mathsf{p}_{\mathsf{p}_v} = \mathsf{c}_v) \end{cases} \tag{35}$$

A.5.2 Spanning Tree Construction without Distance to the Root Maintenance

Proof of Theorem 5 Let $\overline{\mathsf{n}}(v)$ denote the color of v's neighbor with the maximum degree, and in case there are several such neighbors, the one with maximum color.

$$\overline{\mathsf{n}}(v) = \{\mathsf{c}_u : u \in N(v) \wedge \delta_u = \max\{\delta_w : w \in N(v)\} \wedge \mathsf{c}_u = \max\{\mathsf{c}_w : w \in N \wedge (\delta_w = \delta_u)\}\} \tag{36}$$

Lemma 10 (Invariant)**.** *For every node $v \in V$ such that $\mathsf{p}_v \neq \varnothing$ and $\mathtt{Ch}(v) \neq \varnothing$, $\overline{\mathsf{n}}(v) \in \{\mathsf{p}_v \cup \mathtt{Ch}(v)\} \neq \varnothing$ remains true.*

Proof. Proof by induction

Basis case: When a node v starts the algorithm it chooses for a parent the node with the maximum degree, if there exist more than one it chooses the one with the maximum identifier among the ones with the maximum degree. So if a node picks a parent u at the first execution of the algorithm it takes $u = \overline{\mathsf{n}}(v)$ so for theses nodes the invariant is preserved. For the nodes v which are a maximum local. We denote by u the node $\overline{\mathsf{n}}(v)$. Suppose that at the first execution of u, u choses the node w as a parent, that means $\mathsf{R}_w > \mathsf{R}_u$ and $\mathsf{R}_w > \mathsf{R}_v$. So after this execution $\mathsf{R}_u = \mathsf{R}_w$, so now v can choose u as a parent and the invariant is preserved for node v. So after one execution of the algorithm for all the nodes the invariant is preserved.

Assumption: Assume true that after t steps of execution, the algorithm preserved the invariant.

Inductive step: Let us consider a node v, by the assumption we have

$$\overline{\mathsf{n}}(v) \cap \{\{\mathsf{p}_v\} \cup \mathtt{Ch}(v)\} \neq \varnothing$$

The node v cannot change its children in can only change its parent, and only if v its a root (see rules $\mathbb{R}_{\texttt{Merge}}$ and $\mathbb{R}_{\texttt{ReRoot}}$ of Algorithm 3). So for $v\ \overline{\mathsf{n}}(v) \in \texttt{Ch}(v)$, the rule $\mathbb{R}_{\texttt{Merge}}$ assigns as a parent a new neighbor u ($u \notin \texttt{Ch}(v)$) so the invariant is preserved. The rule $\mathbb{R}_{\texttt{ReRoot}}$ assigns as a parent of v a child of v so the invariant is preserved.

□

Lemma 11. *The descendants u of v with $\mathsf{new}_u = \varnothing$ have $\mathsf{R}_u \leqslant \mathsf{R}_v$.*

Proof. Proof by induction on the value R_u with u descendants of v

Base case: Each node $v \in V$ with $\mathsf{p}_v = \varnothing$ and $\texttt{Ch}(v) = \varnothing$ has $\mathsf{R}_v = (\deg_v, \mathsf{c}_v, \text{ID}_v)$ and $\mathsf{new}_v = \varnothing$, otherwise an error is detected. A node v takes a parent iff there exists a neighbor w of v such that $\mathsf{R}_w > \mathsf{R}_v$, and in this case v maintains its variable $\mathsf{new}_w = \varnothing$, so the claim is satisfied (see Rule $\mathbb{R}_{\texttt{Merge}}$).

Assumption: Assume that there exists a configuration γ where for every node $v \in V$, all the descendants u of v with $\mathsf{new}_u = \varnothing$ have $\mathsf{R}_u \leqslant \mathsf{R}_v$.

Inductive step: We consider Configuration $\gamma + 1$. For a node v and every descendants u, the assumption gives the property that if $\mathsf{new}_u = \varnothing$, then $\mathsf{R}_u \leqslant \mathsf{R}_v$. Let us now consider the case where there exists a neighbor w of u with $\mathsf{R}_u < \mathsf{R}_w$.

- If w is the parent of u, R_u takes the value of R_w (see rule $\mathbb{R}_{\texttt{Update}}$). By the induction assumption, we have $\mathsf{R}_w \leqslant \mathsf{R}_v$ (as a parent of u, w is also a descendant of v). So, R_u remains inferior or equal to R_v.
- By the induction assumption, if $\mathsf{R}_w > \mathsf{R}_u$, then w cannot be a descendant of u.
- If w is not in the same subtree of u, u cannot change its parent because u is not a root (see rule $\mathbb{R}_{\texttt{Merge}}$). So u changes its R_u to R_w, but it sets $\mathsf{new}_u = w$ (see rule $\mathbb{R}_{\texttt{Path}}$).

Now, if there exists a neighbor w of u such that $\mathsf{R}_w = \mathsf{R}_u$, then to execute rule $\mathbb{R}_{\texttt{ReRoot}}$, u must be a root. We obtain a contradiction with our assumption that u is a descendant of v.

To conclude, if u is the descendant of v in configuration γ and it remains a descendant of v at configuration $\gamma + 1$, and the value of new_u remains empty, then $\mathsf{R}_u \leqslant \mathsf{R}_v$ in configuration $\gamma + 1$.

□

Lemma 12. *If there exists an acyclic spanning structure in Configuration γ, then any execution of a rule maintains an acyclic spanning structure in Configuration $\gamma + 1$.* [2]

Proof. Proof by induction on the size of the acyclique spanning structure.

Basis case: By contradiction: Remark that, thanks to Algorithm **C-Color**, there exist a total order between the neighbors of a node. Let us consider three neighbor nodes $a, b, c \in V$ such that in Configuration γ, a, b and c have no relatives. Then all three nodes are enabled by rule $\mathbb{R}_{\texttt{Merge}}$. Let us suppose for the purpose of contradiction that in Configuration $\gamma + 1$ a cycle exists. More precisely: $\mathsf{p}_a = b$, $\mathsf{p}_b = c$ and $\mathsf{p}_c = a$, to achieve that :

1. a must choose b as a parent, for that $\mathsf{R}_b > \mathsf{R}_c$

[2] When a node v has $\mathsf{p}_v \in \texttt{Ch}(v)$, we delete p_v (see rule $\mathbb{R}_{\texttt{Del}}$), so we do not consider this case as a cycle.

2. b must choose c as a parent, for that $\mathsf{R}_c > \mathsf{R}_a$
3. c must choose a as a parent, for that $\mathsf{R}_a > \mathsf{R}_b$

We obtain a contradiction between (1),(2) and (3).

Assumption: Assume true that in configuration γ there exists an acyclic spanning structure.

Inductive step: Let us consider a node $v \in V$, a node v takes a new parent only in two cases, and in both case v must be a root.

Let us consider first rule $\mathbb{R}_{\texttt{Merge}}$, let u be the neighbor of v with $\mathsf{R}_v < \mathsf{R}_u$ and $\mathsf{new}_u = \varnothing$. By Lemma 11, u is not a descendant of v, so if v takes u as a new parent, an acyclic spanning structure is preserved. Now, we consider rule $\mathbb{R}_{\texttt{ReRoot}}$. Let u be the neighbor of v such that $\mathsf{R}_u = \mathsf{R}_v$ and $\mathsf{new}_v = u$. In this case, u is either a child of v with $\mathsf{new}_u \neq \varnothing$, or v is not a child of v with $\mathsf{new}_u = \varnothing$. If v is a child of v, v takes u as a parent. Remark that the first action of u is to delete its parent (see rule $\mathbb{R}_{\texttt{Del}}$, and consider the fact that all other rules require $\mathsf{p}_v \notin \texttt{Ch}(v)$), so we do not consider this case as a cycle. If v is not a child of v with $\mathsf{new}_u = \varnothing$, by Lemma 11 u is not a descendant of v. Now, when v takes u as a parent, this action maintains an acyclic spanning structure.

To conclude, Configuration $\gamma + 1$ maintains a acyclic spanning structure.

□

Lemma 13. *If v is a node such that $\mathsf{R}_v = \mathsf{R}_r$, and every ancestor of v (and v itself) have $\mathsf{new} = \varnothing$. Then r is an ancestor of v, or v itself.*

Proof. Suppose for the purpose of contradiction that $r \neq v$, and r is not an ancestor of v. By Lemma 12, v is an element of a sub spanning tree $\texttt{T}$. Let w be the oldest ancestor of v such that $\mathsf{new}_w = \varnothing$. By hypothesis, every ancestor z of v (including w) has $\mathsf{R}_z \neq \mathsf{R}_r$. By Lemma 11, we have $\mathsf{R}_v \leqslant \mathsf{R}_z$, which contradicts $\mathsf{R}_v = \mathsf{R}_r$. □

Lemma 14. *Executing Algorithm* **ST** *constructs a spanning tree rooted in node with the maximum degree, maximum color, maximum identifier, assuming the state model, and a distributed unfair scheduler.*

Proof. Let $\psi : \Gamma \times V \to \mathbb{N}$ be the function defined by:

$$\psi(\gamma, v) = \big((\deg_\ell - \delta_v) + (\mathsf{c}_\ell - \mathsf{c}_v) + (\text{ID}_\ell - \text{ID}_v)\big)$$

where ℓ is the node with $\mathsf{R}_\ell > \mathsf{R}_v$ with $v \in V \backslash \{\ell\}$. Now, let $\phi : \Gamma \times V \to \mathbb{N}$ be the function defined by:

$$\phi(\gamma, v) = \begin{cases} 2 & \text{if } \mathsf{new}_v \neq \varnothing \\ 1 & \text{if } \mathsf{p}_{\mathsf{p}_v} = \mathsf{c}_v \\ 0 & \text{otherwise} \end{cases}$$

Remark that a node v cannot have $\mathsf{new}_v \neq \varnothing$ and $\mathsf{p}_v \in \texttt{Ch}(v)$. Otherwise an error is detected through predicate $\text{ErST}(v)$.

Let $\Psi : \Gamma \to \mathbb{N}$ be the potential function defined by:

$$\Psi(\gamma) = \Big(\sum_{v \in V} \psi(\gamma, v), \sum_{v \in V} \phi(\gamma, v) \Big).$$

Let γ be a configuration such that $\Psi(\gamma) > 0$, and let v be a node in V such that v is enabled by a rule of Algorithm **ST**. If v executes rules $\mathbb{R}_{\texttt{Update}}$, $\mathbb{R}_{\texttt{Merge}}$, or $\mathbb{R}_{\texttt{Path}}$, then R_v increases and we obtain $\psi(\gamma', v) < \psi(\gamma, v)$. Now, if v executes rule $\mathbb{R}_{\texttt{ReRoot}}$, this implies new_v is not empty. After execution of $\mathbb{R}_{\texttt{ReRoot}}$, new_v become empty, so $\phi(\gamma, v)$ decreases by one. Finally, if v executes rule $\mathbb{R}_{\texttt{Del}}$, it implies that v had $\mathsf{p}_{\mathsf{p}_v} = \mathsf{c}_v$, and now $\mathsf{p}_v = \varnothing$. As a result, $\phi(\gamma', v) = \phi(\gamma, v) - 1 = 0$.

Therefore, we obtain $\Psi(\gamma') < \Psi(\gamma)$. By Lemmas 12 and 13 we obtain the property that when $\Psi(\gamma) = 0$, a spanning tree rooted in ℓ is constructed. □

Lemma 15. *Algorithm* **ST** *converges in* $O(\Delta n^3)$ *steps.*

Proof. Direct by the potential function $\Psi(\gamma)$. □

Lemma 16. *Algorithm* **C-ST** *converges in* $O(\Delta n^3 \log n)$ *steps.*

The proof of Theorem 6 mimics the proof of Theorem 5.

A.6 Corrects of the algorithm of Leader Election

Proof of Theorem 1. We first need to show that the number of activations of rules of algorithm **C-ST** are bounded if there exist nodes enabled by **C-Color**, **Freeze** or **C-Break**. Let us consider a subset of the nodes A enabled for at least one of these algorithms, and by S the nodes enabled by rule **C-ST**. The nodes in S belong to some spanning trees (possibly only one), otherwise at least one of rules of **Freeze** or **C-Break** would be enabled. So, there exist a node in S that is enabled by algorithm **C-ST**. Algorithm **C-ST** is talkative, but it runs by waves, and its waves require that all neighbors of a node v have the same R at each phase. As we consider connected graphs only, there exists at least one node v in S with a neighbor u in A. Then, there exists a configuration γ' where the rules of **C-ST** are not enabled, because u cannot have the same R at each phase (since u is not enabled by rules of **C-ST**). So, only Algorithms **C-Color**, **Freeze**, and **C-Break** may now be scheduled for execution, as they have higher priority. As they are silent and operate under an unfair distributed deamon, we obtain convergence.

Let us now consider a configuration γ where no node are enabled for Algorithm **C-Color**, **Freeze**, and **C-Break**. Yet, there exists a node enabled by Algorithm **C-ST**. Thanks to Theorem 6, we obtain a spanning tree rooted in the node with the maximum degree, maximum color, and maximum identifier. As a consequence, only the root r has $\ell_r = true$ and every other node $v \in V \backslash \{r\}$ has $\ell_v = false$. □